DOCUMENT RESUME

ED 077 682 SE 016 093

AUTHOR Goldin, Gerald A.; Luger, George F.
TITLE Artificial Intelligence Models for Human Problem-Solving.
PUB DATE Feb 73
NOTE 77p.; Paper presented at the annual meeting of the American Educational Research Association, New Orleans, Louisiana, February 1973

EDRS PRICE MF-$0.65 HC-$3.29
DESCRIPTORS Cognitive Development; Learning; *Learning Theories; Mathematical Models; *Mathematics Education; *Problem Solving; *Research; Thought Processes

ABSTRACT

A theory that there is a correspondence between Piagetian conservation operations and groups of symmetry transformations, and that these symmetry transformations may be used in explaining human problem solving behaviors, is developed in this paper. Current research in artificial intelligence is briefly reviewed, then details of the symmetry transformation theory are given along with examples of its application. The Tower of Hanoi Problem is extensively analyzed as an illustration of the theory. (DT)

FILMED FROM BEST AVAILABLE COPY

ED 077682

U.S. DEPARTMENT OF HEALTH, EDUCATION & WELFARE
NATIONAL INSTITUTE OF EDUCATION
THIS DOCUMENT HAS BEEN REPRODUCED EXACTLY AS RECEIVED FROM THE PERSON OR ORGANIZATION ORIGINATING IT. POINTS OF VIEW OR OPINIONS STATED DO NOT NECESSARILY REPRESENT OFFICIAL NATIONAL INSTITUTE OF EDUCATION POSITION OR POLICY.

ARTIFICIAL INTELLIGENCE MODELS

FOR

HUMAN PROBLEM-SOLVING

by

Gerald A. Goldin and George F. Luger

Mathematics Education Research

Graduate School of Education University of Pennsylvania

Philadelphia, Pennsylvania 19174

February 1973

SE 016 093

Abstract

Artificial intelligence models are increasingly employed to describe human problem-solving. Here the relationship is developed between such research, and Paigetian or more generally "structuralist" theories of cognition. A fundamental correspondence is suggested between Piagetian conservation operations and groups of symmetry transformations. Acquisition of the ability to treat distinct states of a problem as equivalent when they are symmetrically conjugate, may be a basic process in the development of cognitive structures.

It is further suggested that the decomposition of a problem into subgoals and subproblems may affect problem-solving behavior, even if the infrastructure of subproblems within the main problem is not on the surface apparent.

The state-space representation of a problem, borrowed from artificial intelligence theory, is utilized to define these concepts more precisely and to investigate their consequences. The actual behaviors of subjects solving a problem may be represented by paths through the state-space. Based on the theoretical ideas set forth, hypotheses are suggested predicting certain patterns in such paths--for example the predominance of goal- and subgoal-directed paths, and the presence of congruent paths through isomorphic subproblems.

The Tower of Hanoi problem is used to illustrate the main ideas discussed. The paths through the state-space generated by two individual subjects display the predicted effects of the subproblem decomposition and of the symmetry within the state-space. A natural

distinction emerges between cognitive structures, i.e. conservation operations and symmetries which distinguish the states themselves, and problem-solving strategies for proceeding within the state-space.

An Appendix by one of the authors (GFL) tabulates the behaviors of forty-five adult subjects solving the Tower of Hanoi problem, in a test of the suggested hypotheses.

I. Introduction

A dramatically increasing body of research employs "artificial intelligence" models, or mechanical models, to describe human problem-solving. Some of this research is oriented towards finding the most efficient algorithms or strategies for solving problems with a machine (Arbib, 1969; Banerji, 1969; Nilsson, 1971; Minsky and Papert, 1972), while other research is directed towards simulating or modelling the human being as a problem-solver (Johnson, 1964; Newell and Simon, 1972). The present paper, proceeding in the spirit of Newell and Simon, develops what the authors believe to be a heretofore neglected relationship between artificial intelligence research, and Piagetian or more generally "structuralist" theories of cognition.

Two main theoretical ideas are introduced. First, we assert the fundamental correspondence between Piagetian conservation operations and symmetry transformations. In its most general sense, a symmetry transformation is any operation which carries one state or situation into another in such a way as to leave unchanged important observable features. In the everyday sense of the word "symmetry," these features are geometric; for example, the transformation which changes a particular configuration of objects into its "mirror image" may leave the appearance of the configuration unchanged. However, we shall be interested in syntactic symmetries and symmetries of the underlying structures of mathematical problems, as well as in the more readily apparent geometric symmetries.

Attention is focussed on a subject's ability to treat perceptually distinct states of a problem as equivalent, when such states are related

by virtue of a symmetry transformation. The acquisition of such an ability is frequently essential to the correct solution of a problem, and seems to correspond to the "insight" phenomenon described by the Gestalt psychologists. The present paper suggests that "symmetry acquisition" may actually be as fundamental a process in the development of cognitive structures, as is the acquisition of conservation operations.

The second idea which the authors pursue is that in problem-solving, the subject effectively decomposes a problem into subgoals and subproblems. Such a decomposition may govern a subject's behavior even when he has not consciously directed himself towards a particular subgoal, and despite the fact that the structure of subproblems within the main problems may not on the surface be apparent. Under such an hypothesis, one kind of symmetry which may be explored in a problem is the presence of subproblems having identical or isomorphic structure.

The "state-space representation" of a mathematical puzzle or problem, borrowed from mechanical problem-solving (artificial intelligence) theory, is utilized to define the above concepts more precisely, and to investigate their consequences. The actual behaviors of subjects solving a problem may be represented by paths through the state-space, corresponding to the sequence of steps the subject takes or the moves he makes. The theoretical ideas set forth in this paper lead to the prediction of certain recurrent patterns in such paths--for example, the predominance of goal- and subgoal-directed paths, and the presence of congruent paths through isomorphic subproblems.

Section II is a short review of the necessary background and the most applicable current research in artificial intelligence.

Section III introduces the relationship between symmetry transforma-

tions and conservation operations, and draws an analogy with the physical sciences in order to motivate this relationship. The notion of symmetry is discussed in the state-space of a problem. The example of Piagetian number conservation is examined in detail; examples are also drawn from Tic-Tac-Toe, 2-pile Nim, and a Checkerboard Problem.

In Section IV, we define additional concepts central to the present approach to human problem-solving: subgoals and subproblem state-spaces, isomorphisms and automorphisms in the state-space, and various means by which a subject may "reduce" the state-space diagram in accomplishing the solution of a problem.

The Tower of Hanoi problem, also known as the Tower of Brahma (Gamow, 1947; Gardner, 1959), is introduced in Section V, and its state-space used to illustrate the main ideas so far discussed. The paths through the state-space generated by two individual subjects are displayed. These illustrate the predicted effects of the decomposition into subproblems, and of the presence of symmetries within the state-space. Some possible implications of the present research are suggested in Section VI.

A paper to follow by one of the authors (GFL) tabulates the behaviors of forty-five adult subjects solving the Tower of Hanoi problem, and investigates the validity of the hypotheses presented here.

II. Review of Current Research

Here we shall summarize some of the techniques of mechanical problem-solving, or "artificial intelligence," which have found application to the programming of problem-solving capabilities on the computer. These are the techniques from which we borrow in order to establish a framework for the discussion of human problem-solving. We shall also mention some of the approaches by other authors applying artificial intelligence methods to describe or model the problem-solving behavior of human beings.

A. State-Space Representations and Search Algorithms

Nilsson (1971) defines the "state-space representation" of a problem as the set of distinguishable "situations" or configurations of the problem, together with the permitted "moves" or steps (transitions) from one problem situation to another. Thus a problem consists of an initial state, together with all of the states which may be obtained from the initial state by successive moves. One or more of these successor states is classified as a goal state.

For example, the problem might be to prove a given theorem in symbolic logic. The initial state would be the set of premises of the theorem, a collection of well-formed formulas. A state of the problem would be any collection of well-formed formulas which could be obtained from the initial set by successive application of the rules of logic. The application of a simple rule of logic to add a new well-formed formula to those previously obtained would constitute a permitted step or transition from one state to the next. A goal state would be any set

of well-formed formulas which included the desired conclusion of the theorem.

A generalization of the concept of a state-space representation for a problem is the analogous structure for an N-player game. A problem may then be considered a 1-player game; examples of 2-player games are chess, checkers, and tic-tac-toe. A state is now defined as any configuration of the game, with the additional information as to which player has the move included in the description of the game configuration.[1] For example, a state in chess is any legally reachable position of the pieces, together with the information that White (or Black) has the move. The legal moves of the game determine the transitions from state to state. In the game tree, the opposing players typically have different goal states or disjoint sets of goal states. The game tree for 2-pile Nim is depicted in Figure 1.

One goal of artificial intelligence research has been to program high-speed computers to solve problems in logic, to play games such as chess and checkers, or to make decisions based on available information in arbitrarily specified situations to obtain the most favorable probable outcome. Thus an entire branch of this research is devoted to obtaining efficient search algorithms by means of which the computer can "look ahead" in the state-space or game tree, or "foresee" the possible outcomes following a particular choice. Nilsson discusses "breadth-first" and "depth-first" algorithms for searching within the state-space or game tree, as well as strategies that combine features of these two approaches.

---

[1] Exceptions are games in which players move simultaneously without knowledge of opponents' moves.

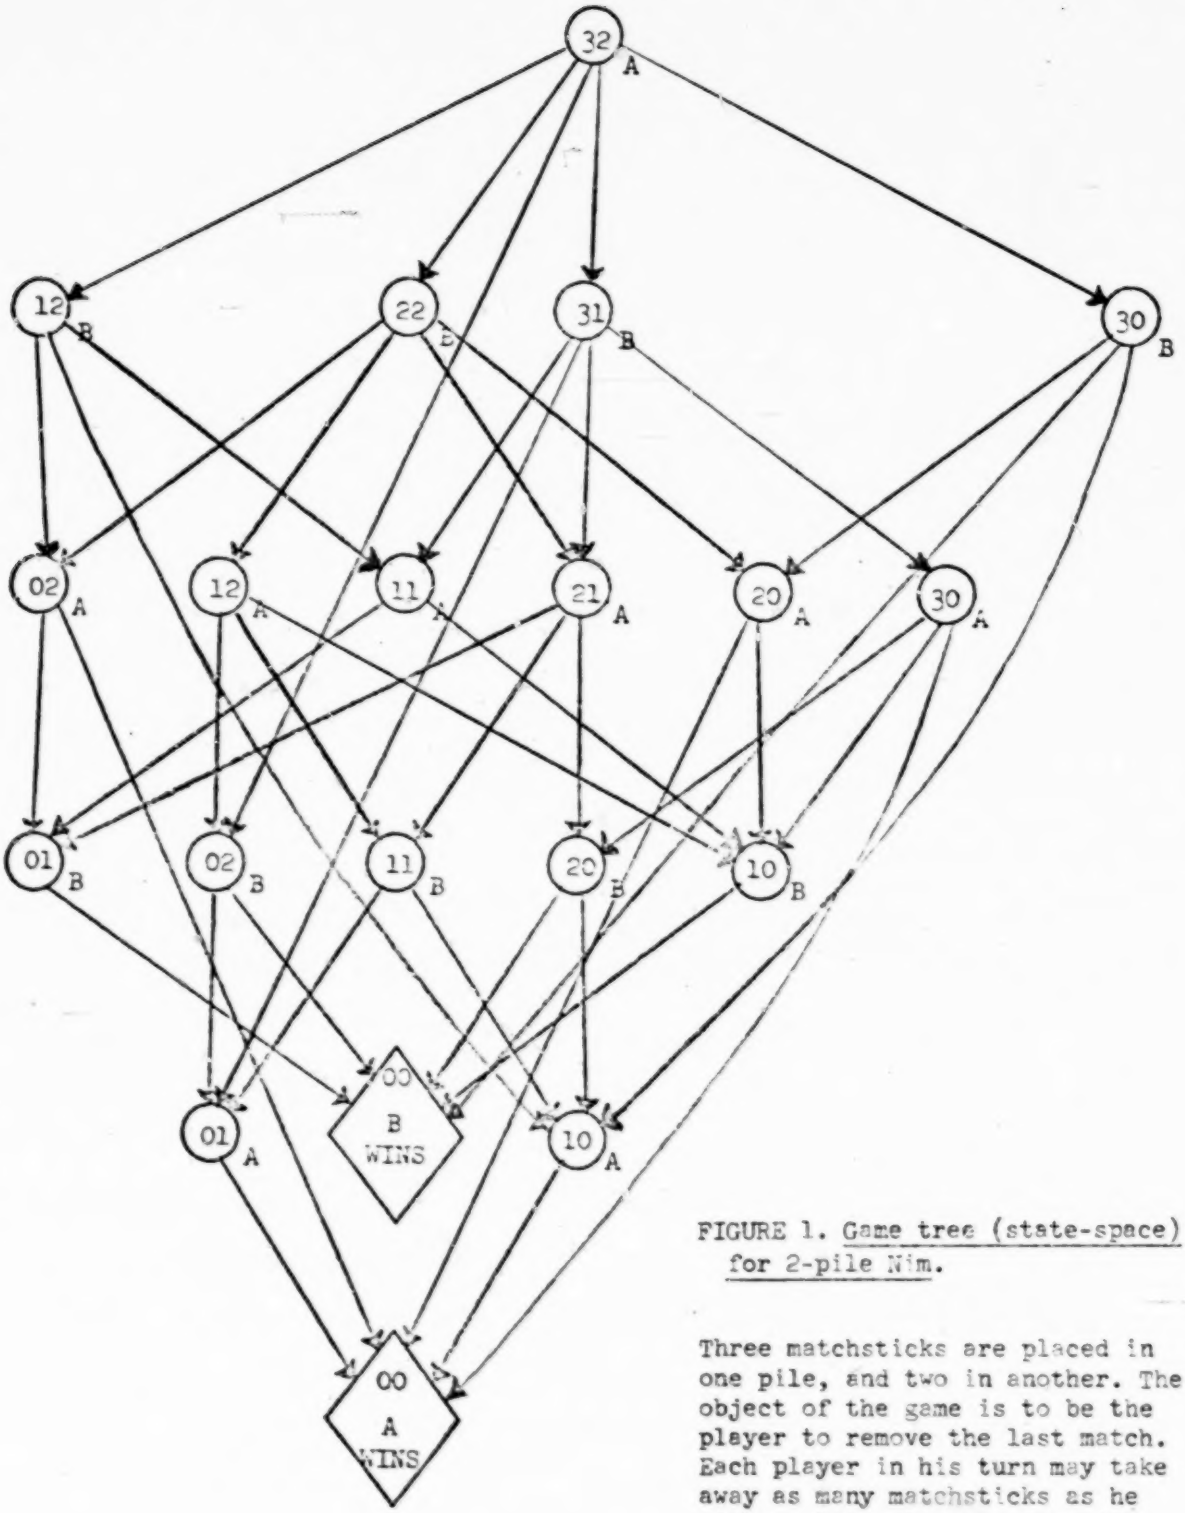

FIGURE 1. Game tree (state-space) for 2-pile Nim.

Three matchsticks are placed in one pile, and two in another. The object of the game is to be the player to remove the last match. Each player in his turn may take away as many matchsticks as he wishes, but only from one pile. Each state is designated by a pair of numbers representing the matches remaining in the respective piles, and by a subscript denoting the player who has the move.

Since for most problems or games the number of possible branches rapidly becomes astronomical, the field of choice must somehow be narrowed. In order to avoid searching to the very end of every path, a value may be assigned to each state based on an observable feature of that state, which represents a measure of expectation for future success. An example of this technique is the use of "positional judgment" in chess whereby such features as "control of the center" and "safety of the King" render a position desirable or undesirable. Once criteria for such an evaluation have been established, the search algorithm may be constructed so as to look only n moves into the future, to calculate the evaluation function for the terminal states thus reached, and to make a choice which maximizes the minimum value of all terminal states resulting from that choice (Nilsson, 1971, p. 138). That is, under the assumption that the player's opponent(s) make the best possible moves in all cases, such an algorithm maximizes the expectation of success. A modification of the above "minimax" procedure which further reduces the number of states in the search is to make a selection based on certain pre-specified heuristics of the moves whose continuations are to be investigated.

The "General Problem-Solver" of Newell, Shaw and Simon (1959) embodies a kind of depth-first search algorithm in which the first object of the program is to identify a subgoal state which might eventually lead to solution of the main problem. The subgoal state is chosen to be "less distant" in some suitable sense from the goal state, than is the initial state. When such a subgoal has been identified, control switches to the task of attaining the subgoal, prior to returning to the main problem. This technique is to be applied recursively, until a string of attainable

subgoals has been generated that extends from the problem's initial state to its goal state.

With the General Problem-Solver, Newell, Shaw and Simon come closest in the area of artificial intelligence to taking the position that utilization of the subgoal and subproblem structure of a problem is fundamental to efficient problem-solving. In the present paper, the authors hypothesize that human problem-solving is demonstrably governed by identifiable subproblems and subgoals within the state-space.

The geometry theorem proving machine of Gelernter (1959, 1960) utilizes the "syntactic symmetries" of a problem to facilitate the search within the state-space. When the program has succeeded in reaching a particular state, it generates those states which are syntactically equivalent, in effect equivalent by symmetry, to the state that was reached, thus obviating the necessity of reproducing all the equivalent paths. Such a program is more efficient in cases where symmetry exists. The present paper asserts the fundamental importance of the symmetries of a problem in influencing the human problem-solver's behavior.

B. Artificial Vs. Human Intelligence

The methods that have been mentioned thus far are all directed towards more efficient machine programming of problem-solving capabilities. While these techniques have often been motivated by some introspectively obtained information as to how a human being might solve the same problem, their main purpose has been effective computer programming. Now let us turn our attention to a different goal, namely the application of aritificial intelligence to the examination, understanding, and modelling of human behavior.

One approach taken by artificial intelligence researchers has been to simulate human problem-solving, human information processing, or human perceptual capabilities. Here the criterion for success has not been any claim that the program actually resembles the way people think, but rather its success in generating human-like behaviors. It is of course impossible to do justice to these programs in a brief review. Perceptual problems such as mechanical procedures for interpretation of depth in two-dimensional scenes (Guzman, 1968) are included among the investigations reported by Minsky and Papert (1972). They also discuss Piaget's conservation experiments from the standpoint of the acquisition of descriptive and deductive procedures. Progress has been made towards the machine interpretation of natural languages (Winograd, 1971). Many efforts along these lines, however, are subject to the limitation that the programming methods employed do not lend themselves to further generalization or extension.

In a different approach from that of trying to simulate human behavior with the computer, Newell and Simon (1972) propose a comprehensive model for the human problem solver as an information-processing system. They introduce a "problem space" to represent the task environment within the information processing system; then they postulate that human problem-solving takes place by means of a search in such a space.

According to Newell and Simon, a problem space consists of:

"1. A set of elements, U, which are symbol structures, each representing a state of knowledge about the task.

2. A set of operators, Q, which are information processes, each producing new states of knowledge from existing states of knowledge.

3. An initial state of knowledge, u_o, which is the knowledge about the task that the problem solver has at the start of problem solving.

4. A problem, which is posed by specifying a set of final, desired states G, to be reached by applying operators from Q.

5. The total knowledge available to a problem solver when he is in a given knowledge state, which includes ... :

   (a) Temporary dynamic information ...
   (b) The knowledge state itself ...
   (c) Access information (to memory) ...
   (d) Path information about how a given knowledge state was arrived at ...
   (e) Access information to other knowledge states that have been reached previously ...
   (f) Reference information ..."

Newell and Simon seek to model observed human performance in problem-solving tasks such as cryptarithmetic, logic theorem-proving, and chess.

The first four components of Newell and Simon's definition of the problem space correspond to taking the state-space representation of the problem as described by Nilsson, interpreting a "state of the problem" as a "state of knowledge about the problem," and incorporating the whole into the information processing system. Rather than combine the problem and the problem solver into one system as Newell and Simon have done, the authors of the present paper prefer to regard them as two separate but interacting systems. Other researchers (Carr, unpublished; Menzel, 1970) also favor the preservation of the distinction between problem and problem solver, utilizing a feedback loop and various decomposition and extension theorems (Arbib, 1969) to generate the successive states of each system. Newell and Simon obtain what they call the "problem behavior graph" of a subject in the "external problem space"--this is the analogue in their model of the paths within the state-space representation which are to be the main objects of study in the present paper.

In a less general context E. M. Johnson (1964) proposes an information processing model simulating the observed behaviors of subjects solving concept formation problems (Bruner, Goodnow and Austin, 1956).

III. Conservation Operations and Symmetries

A. Background

The correspondence between conservation laws and symmetries of nature is a well-known concept in modern physics. For example, conservation of momentum derives from the invariance of physical interactions under spatial translations, conservation of angular momentum from rotational invariance, and conservation of energy from invariance under time translations (Feynman, 1965; Wigner, 1964).

The fact long remained unobserved that such a correspondence existed as a general principle. In some cases physicists became aware of and successfully expressed a conservation law prior to understanding that the law actually derived from a known symmetry of the physical world--for example, in the cases of conservation of momentum, angular momemtum, and energy. In other instances the symmetry was well-known, and physicists proceeded to define an observable whose conservation followed automatically from the fact of obedience to the symmetry. Thus conservation of parity, for instance, follows from the supposed invariance of physical interactions under spatial reflection. Such newly defined observables proved immeasurably useful when it was learned that on a sub-atomic level, symmetries such as spatial reflection which had heretofore been taken for granted were subject to violation, and non-conservation occurred.

Finally there were some well-known conservation laws, based on which previously unknown symmetries could be defined. Thus conservation of electric charge can be interpreted as a consequence of invariance

under rotations in an abstract mathematical space.[2]

It is now understood that the pairing of a conservation law with a symmetry in physics may be regarded as a mathematical rather than an empirical relationship, which follows from the mathematical theory of Lie groups. This relationship asserts that to every set of observables corresponds a certain algebra of observables; and to every such algebra corresponds a group. If the values of the observables are conserved, i.e., unchanged as the system develops in time, then it turns out that the group elements describe physical symmetry transformations of the system.

B. The Structuralist Methodology

The group is the paradigm in mathematics of the methodology which has been termed "structuralist" (Piaget, 1970; Lane, 1970). A group is a set, closed under an associative binary operation, possessing an identity element, and in which each element has a corresponding inverse.

The set of symmetry transformations of a system always forms a group. Any pair of symmetry transformations may be performed successively to generate a third symmetry transformation, defining an associative binary operation. The identity transformation is always included as a symmetry by convention, and to every symmetry transformation corresponds the inverse transformation which returns the system to its initial configuration (Wigner, 1959).

The structuralist methodology has been applied to fields of study as diverse as anthropology (Lévi-Strauss, 1963, 1969), linguistics (Harris, 1951), and psychology (Piaget and Inhelder, 1969), as well as

[2] The space is known to physicists as isotopic spin-space (Eisenbud and Wigner, 1958).

to mathematics (Bourbaki, var.). According to Piaget (1970) a structure in the most general sense is a system or set within which certain relations (or operations) have been defined, embodying the concepts of wholeness, transformation, and self-regulation. For example, a system of kinship constitutes a structure in anthropology as does a group in mathematics.

In Piagetian developmental psychology, the conservation operations --conservation of number, volume, quantity, etc.--are the transformations which govern the cognitive structures assumed to underlie an individual's behavior (Ginsburg and Opper, 1969). The acquisition of these conservation operations by children defines sequential stages in their cognitive development.

In view of the parallel fundamental roles played by group structures in mathematics and the aforementioned cognitive structures in developmental psychology, it is natural to try to look at the acquisition of Piagetian conservation operations as equivalent to the acquisition of a group of symmetry transformations.

For an observable (such as number, quantity, etc.) to be conserved means in fact that when a given state is somehow transformed into an altered state, the value of the observable is unchanged from its initial value. Of course for the second state to be regarded as different from the first at all, there must be at least one other observable which does change in value under the transformation. Such an observable is not conserved by the transformation.

Given a set of states and a set of relationships among them (for example as discussed in Section II, the permissible moves which take one state of a problem or game into another), a symmetry transformation

may be defined as a one-to-one mapping from the set of states onto itself which leaves invariant the specified relationships among the states. Any collection of such symmetry transformations generates a symmetry group.

Let us say that a given symmetry group G conserves a given set of observables when for every state S in the system, all states which may be obtained from S by applying symmetry operations from G have exactly the same values of the specified observables. We shall also be interested in the maximal symmetry group possessing this property for a given set of observables, that is, every symmetry transformation which preserves the values of the specified observables is to be included in the group.

As an example, consider the rearrangement of n objects on a table or two-dimensional surface depicted in Figure 2. The final configuration of objects (described by the coordinates $\vec{x}_1', \ldots, \vec{x}_n'$) may be obtained from the initial configuration $(\vec{x}_1, \ldots, \vec{x}_n)$ by means of a rearrangement mapping or deformation which appropriately transforms the points in the 2-dimensional plane. Such a rearrangement must be one-to-one (so that two objects do not wind up at the same point) and is taken to be surjective (so as to be invertible). Noting that any two mappings of this kind may be applied successively to yield a third, the set of all such mappings forms a group K. For this example, the collection of states is the set of all possible configurations of n objects in 2-dimensional space, for $n = 0, 1, 2, \ldots$.

To say that "number is conserved" means that when a given state (of say n objects) is transformed into an altered state by moving the objects around (not by adding any or taking any away), then the value of the observable "number" remains unchanged at n. The group K defined

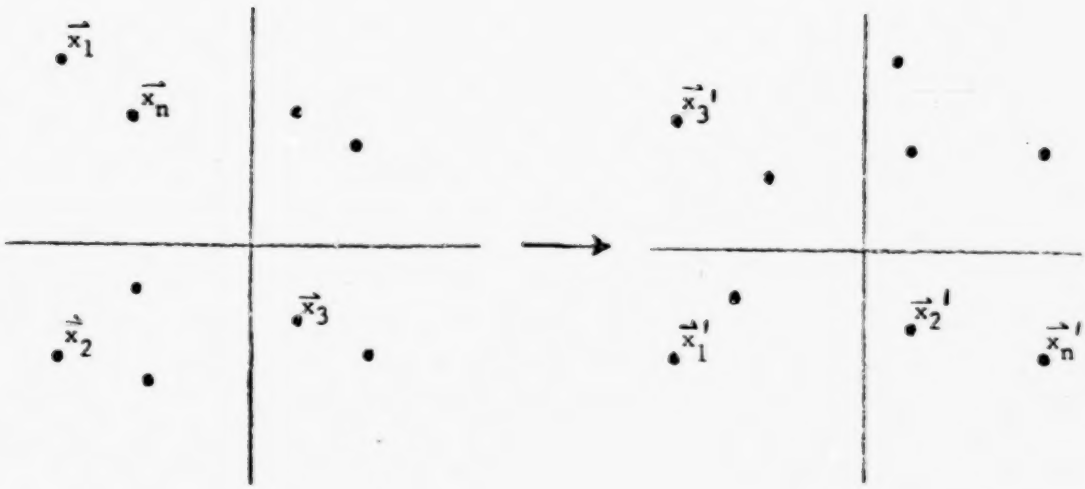

FIGURE 2. Rearrangement of n objects in 2-dimensional space.

The transformation may be implemented by means of a spatial rearrangement mapping or deformation.

above, that is the group of one-to-one surjective mappings from $\mathbb{R}^2$ onto $\mathbb{R}^2$, maps the set of states onto itself in such a way that a state specified by n points continues to be specified by n points after it is transformed, and thus has the same value of the observable "number." It is not difficult to see that K fits our definition of a symmetry group conserving that observable.

What we are saying is that in principle the acquisition of "number conservation," that is the ability to respond that the number of objects remains unchanged when only the positions of the objects have been changed, is logically equivalent to the acquisition of the structure of the symmetry group K, that is, the ability to undo (invert) any rearrangement transformation and to perform any two such transformations successively.

It may well be that stages in the acquisition of a symmetry group structure actually correspond to the acquisition of particular subgroups. For example, a child at some time might recognize that the number of objects is conserved when a configuration is merely translated a certain distance in space, without its being spread out or otherwise rearranged. If this were to occur, we would be able to say that the subgroup of K composed of all translations had been acquired as a symmetry structure.

The symmetry group for the above example of number conservation is relatively complicated to define. Furthermore its elements are only "symmetries" in a rather formal sense, namely that different configurations of n objects may successfully be treated as equivalent for some purposes. Having argued for the reformulation of conservation operations in terms of symmetry groups, the next step is to cite examples of systems in which the symmetries are more familiar, but where the

identification of conserved quantities is more cumbersome. Such examples are considered in the next sub-section. Unlike the case of conservation of number, many examples drawn from problem-solving turn out to be easier to describe in terms of symmetry groups than in terms of quantities conserved by the transformations in those groups.

The above correspondence between symmetry groups and conserved quantities is analogous but not identical to the correspondence in physics between symmetries and conservation laws outlined in Section III.A. The major difference is that for a physical system the time-development operator plays a special role. A physical state evolves in time uniquely according to the dynamics of the system; and conservation of an observable quantity means specifically that the quantity is unchanged by the time-development. On the other hand in attempting to describe problem-solving behaviors, one must allow for a choice of possible moves. While there may be many equivalent choices in accordance with the symmetry that is present, the problem-solver must nevertheless make only one choice, and cannot make all of them simultaneously. Therefore the <u>actual</u> time-development will often be asymmetrical.

Thus for a physical system conservation is defined with respect to the time-development operator only, and states that are conjugate by symmetry remain so as they develop in time. In the state-space of a problem or game, however, we look at the group of <u>all</u> transformations which conserve the values of the specified observables, a group which most often does not include any transformation which effects the actual time-development.

C. Symmetries in the State-Space Representation

The characterization of the "states" of a system is the essential first step in cybernetics and virtually all artificial intelligence research. Specification of two states as distinct means that they differ in the value of one or more observables, i.e., in a quantity which is not conserved by operations connecting the two states.[3] Alternatively, characterization of two states as equivalent for the purpose of solving a given problem means that they have the same values for the observables relevant to that problem; in fact they are equivalent modulo a symmetry transformation which leaves invariant the values of those observables and preserves the relationships among the states.

1. Tic-Tac-Toe

For a simple example, consider the state-space representation of the game Tic-Tac-Toe. There are nine distinguishable states which can be reached by the first move of the first player. However, modulo the rotation or reflection symmetry, only three distinguishable states exist (Figure 3). In constructing the state-space representation for Tic-Tac-Toe, one could choose to represent all the distinguishable states of the system, thus obtaining a very large state-space; or one could use the much smaller state-space obtained by regarding those states related by symmetry as equivalent. Choice of the smaller state-space corresponds to "reduction" of the state-space diagram modulo the symmetry transformations.

Thus in originally specifying the state-space representation for

---

[3] It should be remarked that the value of an observable which characterizes a state need not be a numerical value. For example, the observable "color" might have the value red, green, or yellow, in specifying the states of a traffic directional system.

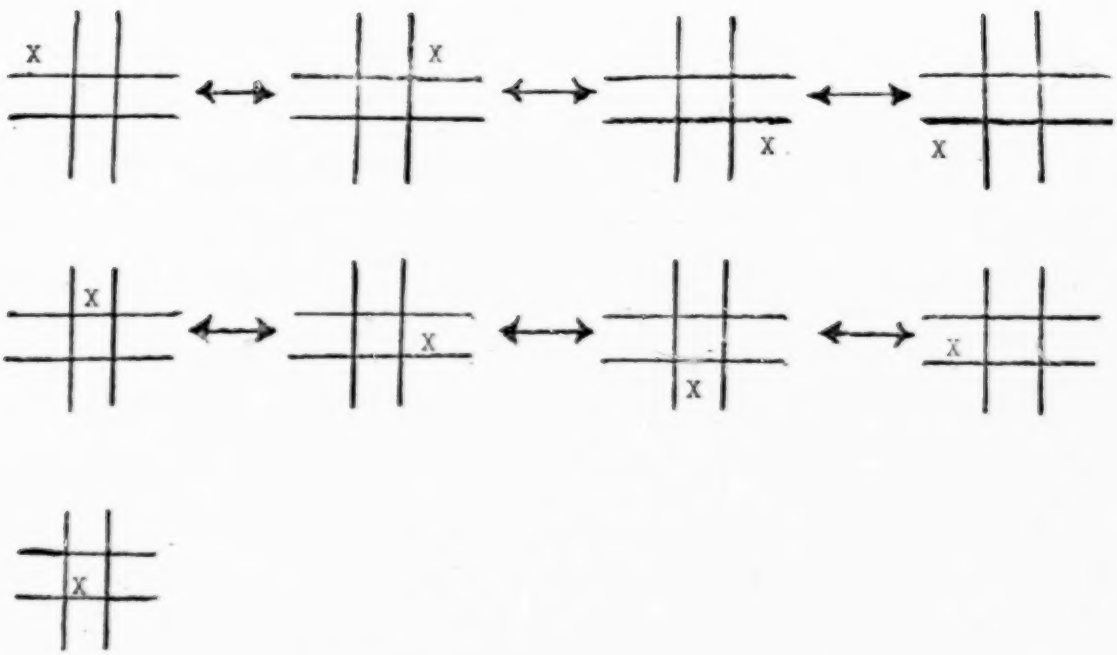

FIGURE 3. Tic-Tac-Toe states equivalent by symmetry.

a problem or game, one must make a choice: to incorporate or to neglect the available symmetries. From the standpoint of efficient problem-solving, the machine programmer will try to incorporate as much of the symmetry as possible. However, in studying actual human problem-solving, we must take into account the possibility that the subject solving a problem or playing a game does not necessarily perceive all of the symmetry which is in fact present. Therefore if we are to map his behaviors faithfully, we must begin with the expanded state-space representation of the problem or game.

Tic-Tac-Toe provides an example of a game in which the rotation and reflection symmetry is easily recognized, but the corresponding conserved quantities are cumbersome to define. For example, one such quantity would be the number of X's in corner squares; another, the number of 0's in corner squares. These numbers are unchanged by the rotation or reflection operations. In order to characterize the distinguishable states completely, it would still be necessary to separate the situation of two X's (or 0's) in opposite corners from two X's (or 0's) in adjacent corners; this could be done by means of still another numerical observable. Similar observables can be defined which describe the occupancy of the side squares and of the center square. In learning the game of Tic-Tac-Toe, one of the steps might be to realize that if the first player has put an X in the center, the second player can force a draw if he places an 0 in <u>any corner</u>, but loses if he places an 0 in <u>any side</u> square. This formulation clearly incorporates the symmetry in its emphasis on the geometric property of "corner" or "side." It can be restated in terms of the numerical observables defined above.

A game isomorphic to Tic-Tac-Toe may be described as follows. The integers 1, 2, ... , 9 are written on a pad, and the two opposing players take turns, each selecting one of the numbers as his own. Neither player may select a number already taken. The goal is to obtain any three numbers which add up to exactly fifteen. Figure 4 illustrates the isomorphism between this game and Tic-Tac-Toe. A player trying to learn this game would not have available the perception of geometric symmetry which is presented by the Tic-Tac-Toe grid. Unless he had prior familiarity with the magic square, he would have to seek such rules as, "If the first player chooses 5, then the second player must pick an even number in order to avoid losing." Posing the problem in this fashion highlights the search for the relevant observables (those which are important to formulation of a successful strategy), which of course, unbeknownst to the player, are just those observables which are conserved by the Tic-Tac-Toe symmetry--"even numbers selected," "odd numbers excluding 5," etc.

In short, the game of Tic-Tac-Toe illustrates (a) that symmetries may be more convenient than the quantities conserved by those symmetries for formulating the notion of equivalency among states, (b) that symmetries and conserved quantities are however logically interchangeable, and (c) that the "rules of the game" may be reformulated in such a fashion as to make identification of the conserved quantities easier or more convenient than characterization of the symmetries.

2. Symmetry in 2-pile Nim

The complete state-space for the game of 2-pile Nim was depicted in Figure 1. While the initial configuration of 3 matchsticks in the

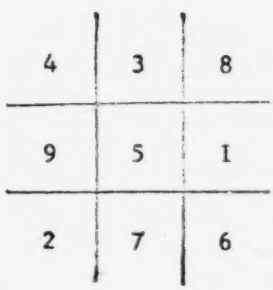

FIGURE 4. Magic Square for the integers 1, 2, ... , 9.

This illustrates the isomorphism between the number selection game described in the text and Tic-Tac-Toe.

first pile and 2 in the second is not symmetrical, it is easy to see that a certain subspace of the state-space is symmetrical, with respect to exchange of the number of matchsticks in the two piles (Figure 5). That is, the state labelled (2,1) with player A to move is equivalent or conjugate to the state labelled (1,2) with player A to move. If player A has a winning strategy in the first situation, he will have an equivalent winning strategy in the second. The "reduced" state-space for a sub-game of 2-pile Nim is indicated in Figure 6.

Thus it need not be the entire game or problem which possesses a symmetry; it is meaningful to discuss the symmetry of a subgame or subproblem.

Finally we remark that while perception of the symmetry in 2-pile Nim does not of necessity mean perception of the winning strategy, it is strongly suggestive. The number of matchsticks in the first pile alone is not conserved by the symmetry operation; but the sum and the difference of the numbers of matchsticks in the two piles are conserved. This suggests that the winning strategy should be formulated solely in terms of these quantities. In fact, the first player can always win by following the rule, "Make both piles equal," or "Make the difference between the numbers of matchsticks in the two piles equal to zero."

3. A Checkerboard Problem

A well-known problem presents the would-be solver with an ordinary checkboard from which two opposing corner squares have been removed, as depicted in Figure 7a. The problem-solver is permitted to cover any two horizontally or vertically adjacent squares at a time with a paper clip. By means of a sequence of such moves, the goal is to cover all of the squares in the original layout.

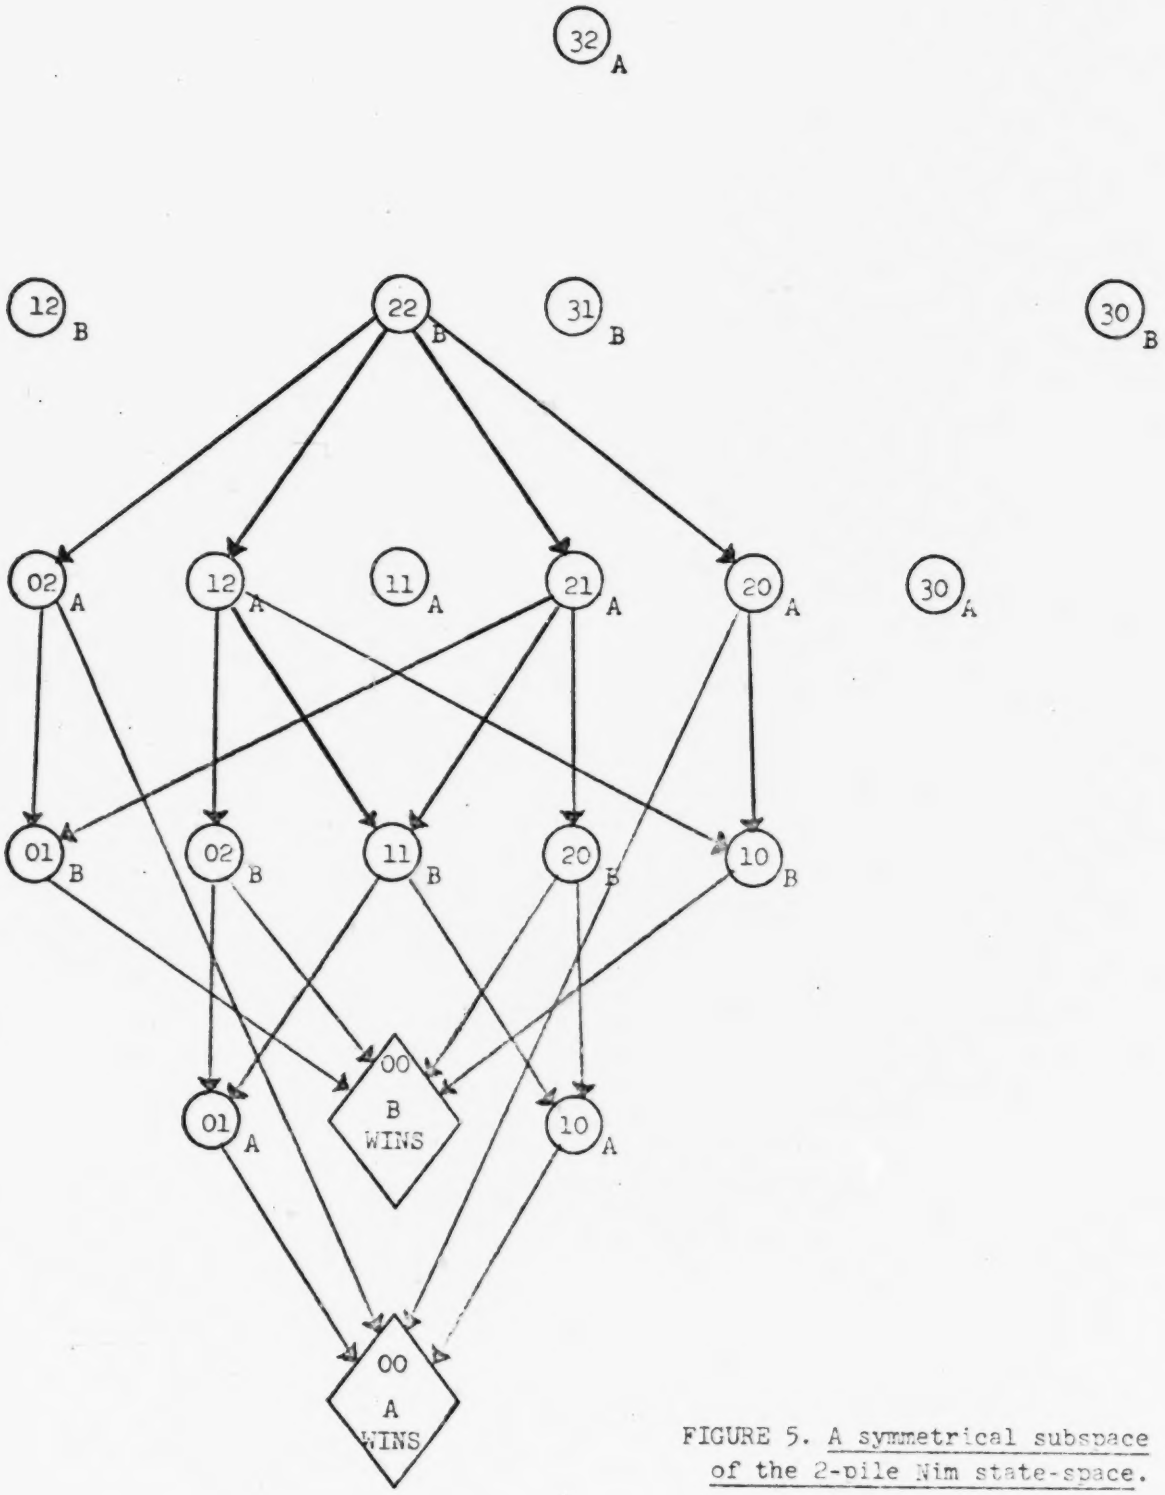

FIGURE 5. A symmetrical subspace of the 2-pile Nim state-space.

Compare with Figure 1.

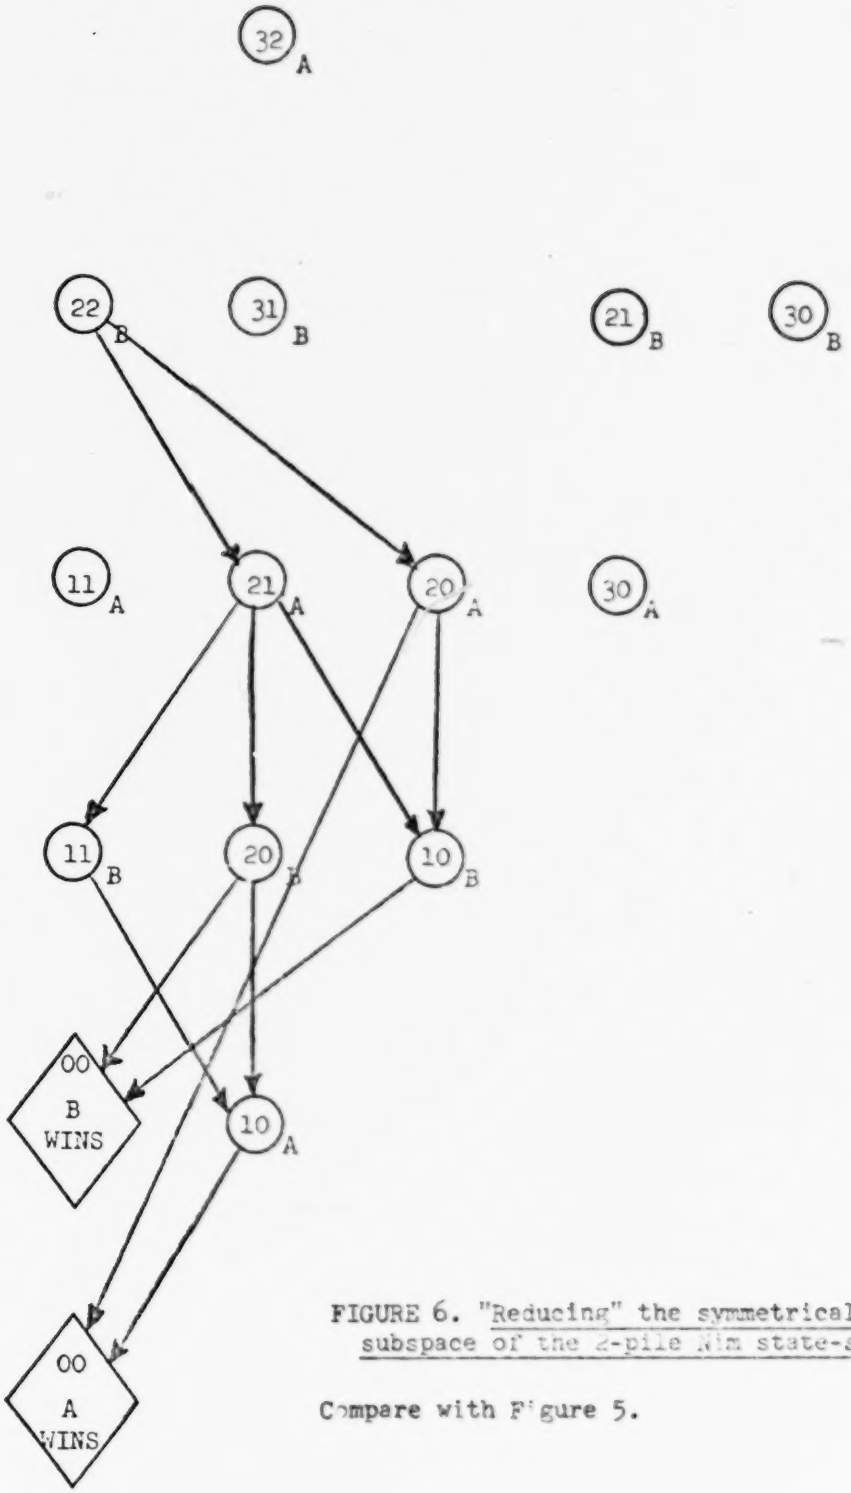

FIGURE 6. "Reducing" the symmetrical subspace of the 2-pile Nim state-space.

Compare with Figure 5.

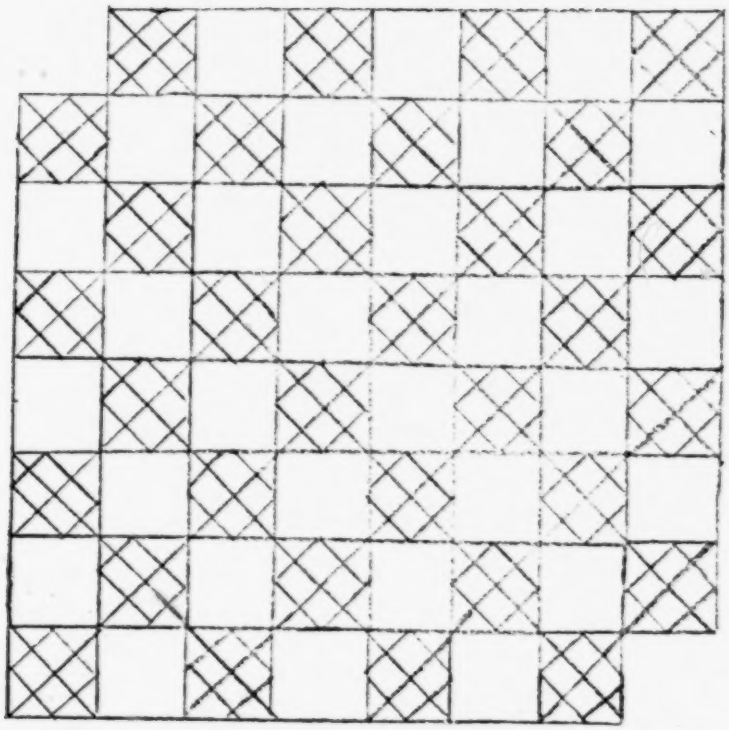

FIGURE 7a. Initial configuration for a checkerboard problem.

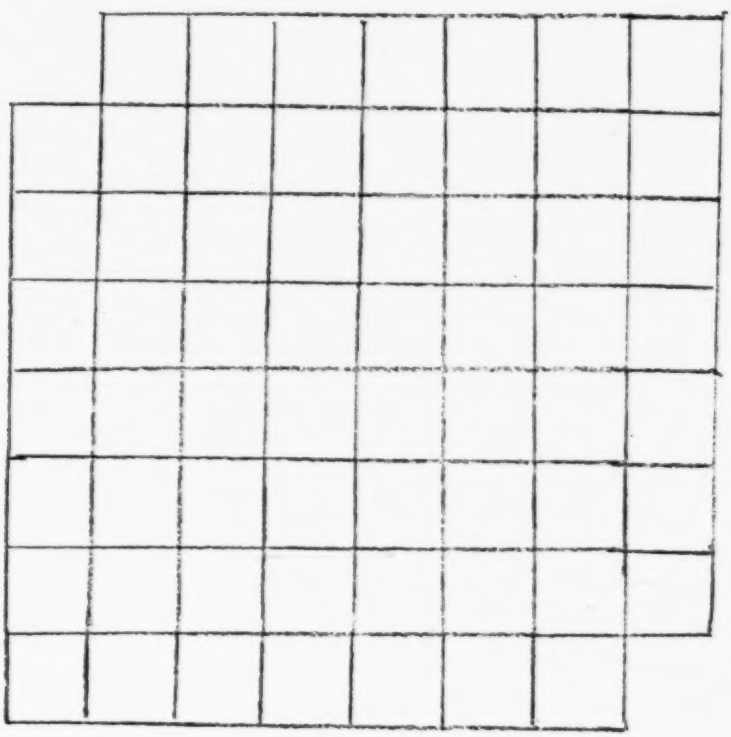

FIGURE 7b. More difficult proposition of the checkerboard problem.

The key to solving this problem is to recognize that the permitted operation of covering two adjacent squares leaves unchanged the quantity "N_b minus N_w," where N_b is the number of white squares remaining, and N_w the number of black squares remaining. However the initial value of this quantity is 2, and its value in the desired goal state is 0. Consequently, the problem as posed is impossible.

Let us construct a state-space for the Checkerboard Problem as follows. A state will be any configuration of black and white squares (smaller in dimension than 8x8) having the checkerboard property that squares of the same color are always diagonal to each other. The state containing no squares, since it is the goal state, must also be included in the state-space; for this reason the state-space must be larger than the collection of states which can be reached by legal moves from the initial state.

Once the fact that the permissible moves conserve the quantity $N_b - N_w$ is noted, then it becomes immediately apparent that the above state-space is composed of disconnected component subspaces, corresponding to different values of $N_b - N_w$; i.e., it is impossible, by legal moves, to travel from one such component subspace to another. The initial state and the goal state simply occur in different components. This problem is unusual in that the transformations which do the conserving (of $N_b - N_w$) are the legal moves themselves, which are thus also "symmetry transformations" in the sense discussed above.

Solution of the checkerboard problem may be made considerably more difficult by proposing it without the shading of the squares (Figure 7b). This removes a perceptual feature which is a clue to identifying the conserved quantity, but which plays no intrinsic role in the state-

ment of the problem itself.

Another famous problem for which the allowed moves themselves conserve a definable quantity is the "15-puzzle" (Figure 8). The state-space contains two disconnected subspaces, which can be derived respectively from the even and odd permutations of the original configuration of number tiles. For example any succession of legal moves which restores the blank to its original position must effect an even permutation of the number tiles. Thus there arises a whole class of "impossible goal" states for any given initial state.

To sum up, several examples drawn from familiar problems or games have been presented in which the structure of the problem manifests itself in the state-space by means of patterns of symmetry. Alternatively it is possible to identify precisely those observable properties of the states which are left unchanged by the symmetry operations. While these two formulations of conservation are logically equivalent, one or the other is frequently more convenient for describing a particular problem situation. It is also meaningful to discuss symmetries that may be present in the state-space of a subproblem of the main problem. Finally the symmetry transformations or conservation laws often provide the key to finding the correct or "elegant" problem solution.

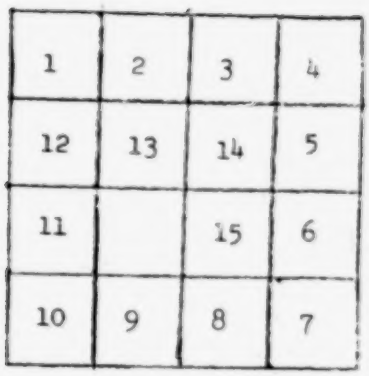

FIGURE 8a. A state of the 15-puzzle.

A legal move is to slide into the blank square one of the number tiles adjacent horizontally or vertically.

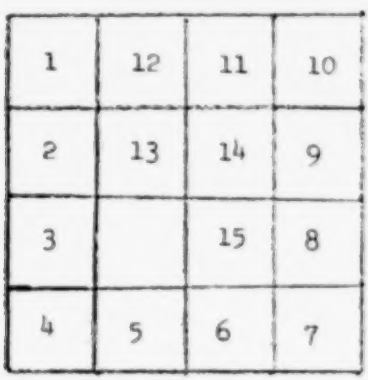

FIGURE 8b. A state of the 15-puzzle impossible to reach from Figure 8a.

IV. Representation of Problem-Solving Behavior

The formal correspondence between a group of symmetry transformations and the observable quantities conserved by those transformations suggests that the acquisition of symmetries may be as fundamental to cognitive development as is the acquisition of conservation operations. We have seen in several problem situations how the presence of symmetry may be represented in the state-space.

A second feature of a problem which is susceptible to study utilizing the state-space is its infrastructure of subproblems. It has been commonly held that an effective problem-solving technique is to establish subproblems or subgoals whose solution or attainment might assist in the conquest of the main problem. Polya (1945) suggests such an approach in discussing his problem-solving "heuristics;" it also forms the basis for Newell, Shaw and Simon's "General Problem-Solver" discussed in Section II.A above, and suggests to Nilsson (1971, p. 80) one way to reduce the state-space. But to establish rigorously the role of such identification of subgoals in human problem-solving behavior remains difficult, and psychologists are still divided even over the assumption of "goal-directedness" (Tolman, 1948; Kimble, 1961, Sec. 13). Characterization of the subproblems of a problem as subspaces of the state-space should assist in investigating the consequences for behavior of a subproblem decomposition by the problem solver.

One may also discuss independently the group of symmetry transformations of a subproblem, as we did in the case of 2-pile Nim (Section III.C). Another kind of "symmetry" whose effects may be explored is the presence in a problem of different subproblems having identical or isomorphic structure.

The above considerations suggest the utility of mapping actual human problem-solving behavior as paths through the state-space representation of the problem. Based on the formal properties of the state-space such as its symmetry and its decomposition into subproblems, hypotheses can be formulated which predict properties of the paths generated. Then the door is open to the development and empirical test of artificial intelligence models for human problem-solving; i.e., general algorithmic or mechanical procedures which would replicate the properties of the paths generated by human problem-solvers. The decision to represent problem-solving behavior as paths through the state-space of the problem is motivated by the desire to make precise what the data is which needs to be "explained" by a theory of human problem-solving. However it does not yet commit us to a particular model or theory.

In practice it may not always be easy to represent behavior in this fashion. The best experimental situation is a problem whose states correspond to different discrete situations of an actual physical device, such as the 15-puzzle or the Tower of Hanoi board. Other available means for recording a subject's behavior as a succession of states entered may include recordings of his oral comments, his written notes, or even his gestures and eye movements (Bartlett, 1958; Newell and Simon, 1972).

Before proceeding with further discussion we shall offer rigorous definitions of the concepts central to the present approach. For completeness some terms are included which have been discussed in earlier sections.

A. Definitions

The state-space of a problem is the set of distinguishable situations

or states of the problem, together with the permitted transitions or moves from one state to another. The problem must specify an initial state, together with one or more goal states.

A problem is impossible if no goal state can be arrived at from the initial state by means of successive transitions.

A subspace of a state-space is a subset of the states, together with the permitted transitions which obtain from one state in the subset to another state in the subset. A subproblem is a subspace of the state-space, having a particular state designated as "initial," and a particular set of states designated as subgoals. For a subproblem it is further required that if the initial state is not the initial state of the main problem, it can be entered from a state outside the subspace; and if a subgoal state is not a goal of the main problem, it can be used to exit from the subspace - i.e., to enter a state of the problem outside of the subproblem. There are often many ways to decompose a particular problem into subproblems, which correspond to different choices of subgoals and corresponding choices of subspaces within the state-space.

Two problems are said to be isomorphic if and only if there is a bijective mapping from the state-space of the first onto the state-space of the second, having the following properties:

1. The initial state of the first problem is mapped onto the initial state of the second.
2. The set of goal states of the first problem is mapped surjectively onto the set of goal states of the second.
3. A transition from one state to another is permitted in the first problem if and only if the corresponding transition is permitted in the second.

Two subproblems of a given problem are said to be isomorphic if they are isomorphic as problems in their own right.

An automorphism of a problem is an isomorphism of the problem onto itself. An automorphism of a problem is also called a symmetry transformation or symmetry automorphism. The set of all of the automorphisms of a problem forms a group (cf. Sec. III.B) under the binary operation of composition or the successive application of two automorphisms. This group is called the symmetry group or automorphism group of the problem.

The states of a problem may be distinguished by virtue of having different discrete values for a set of variables called observables. These observables may or may not be numerical--they may refer to color, position, etc. An observable is said to be conserved by a symmetry transformation (or group of transformations) if and only if for any state the value of that observable is unchanged by the transformation (or group of transformations).

Let S be a state of a problem, and consider the set of all states which can be obtained by applying automorphisms or symmetry transformations from a group G to S. This collection of states is called the orbit of S under the automorphism group G. Two states are said to be congugate modulo the symmetry group G if they are in the same orbit under G.

The orbits within the state-space form mutually disjoint equivalence classes of states. A new and simpler state-space may now be constructed canonically by considering each equivalence class as a state in its own right--or alternatively, by selecting one representative state from each orbit. The state-space thus obtained is said to have been reduced with respect to the symmetry group G, or reduced

modulo G. G may be the full automorphism group of the original state-space, or any subgroup thereof.

A path in the state-space of a problem is a sequence of states $S_1, S_2, \ldots, S_n$ such that for $i = 1, \ldots, n - 1$, the pair $\langle S_i, S_{i+1} \rangle$ represents a permitted transition. A solution path for a problem (or subproblem) is a path in which S_1 is the initial state and S_n is a goal state, with $S_2, \ldots, S_{n-1}$ neither initial nor goal state of the problem (or subproblem).

Two paths within respective isomorphic problems are said to be congruent (modulo the isomorphism) if one path is the image of the other under the isomorphism.

We have seen above that one way to reduce the size of the state-space is with respect to a group of symmetry automorphisms of the problem. A second means of reduction is with respect to the subproblem structure, as follows. The state-space may be described (albeit non-uniquely) as a union of mutually disjoint subspaces, such that for any ordered pair of subspaces, at most one transition exists from a state in the first to a state in the second. Then an entire subspace may be regarded as a state in the reduced state-space, and a transition is permitted from one subspace to another whenever a transition does in fact exist from a state in the one to a state in the other. Each subspace, now a state in the reduced state-space, becomes also a subproblem of the original problem whenever a particular entry state is designated as "initial," and any or all of its exit states are designated as "goals." Then we say that the state-space has been reduced modulo its subproblem decomposition.

Finally one may address the concept of a goal-directed path within a problem or subproblem. Roughly speaking such a path is a solution path which does not "double back" on itself within the state-space, moving consistently "towards" rather than "away from" the goal state in which it terminates. The criteria for defining "doubling back," "distance from the goal state," etc. are for the present to be established in the context of the specific problem under consideration.

The preceding definitions focus on the states and the permitted transitions between them. Let us remark that in any actual problem, at least a portion of the state-space can be generated from the initial state by means of pre-specified rules of procedure (or operators). The operators have verbal descriptions making reference to the values of certain observables for the states on which they act. The goal states may likewise be specified as a class by making reference to the specific properties required of them - again, values of certain observables. These observables may in turn be derived in some complicated fashion from the observables with respect to which the operators are defined.

But the state-space formalism is predominantly concerned with what we may call "problem structure," rather than with alternative means of problem description or with the implications of different embodiments of a single problem.

B. Paths Generated by Problem-Solvers

In problem-solving it may be assumed that the solver acts sequentially upon problem situations (states) to generate successor states, a process which can be described by means of paths through a state-space

constructed by the researcher. It is not suggested that the problem-solver "perceives" the state-space as an entity during problem-solving. The symmetry properties which have been discussed are formal properties of the state-space, which may (as in Tic-Tac-Toe) or may not (as in the Magic Square, Sec. III.C) correspond to geometrical or perceptual properties of the problem readily apparent to the problem-solver.

The state-space description is limited in its immediate applicability to localized problem-solving episodes during which the solver "understands" the rules of procedure, and is able to discriminate between the different values of the perceptual variables which characterize the states. The acquisition of these rules and discriminative abilities prior to the commencement of problem-solving is not explicitly under discussion, although we are intensely interested in such processes during problem-solving as they manifest themselves in altered patterns of behavior.[4]

Nevertheless, some notion of how one intends to proceed from the study of local problem-solving episodes to an understanding of the global process of cognitive change needs to be made explicit, or else one may be foredoomed to conduct a series of merely formal exercises. The present authors view the acquisition of symmetry group structures during problem-solving as an important means of making this transition. The fact of which symmetries are to be incorporated and which are to be neglected ultimately determines which states are to be treated as equivalent and which as distinct. In addition we believe that such

---

[4]Thus Newell and Simon (1972, p. 4) state, "The study is concerned primarily with performance, only a little with learning, and not at all with development or differences related to age."

manifestly global changes as Piaget's cognitive stages can be described in principle in terms of the acquisition of symmetry group structures (Section III.B).

The approach at this stage of the research has been to formulate hypotheses respecting the paths generated by human problem solvers in the state-space of a problem. Such hypotheses are motivated by the formal properties of the state-space under discussion and represent the anticipated effects of the problem structure in shaping problem-solving behavior. The following hypotheses of a more-or-less general nature are suggested.

Hypothesis 1. In solving a problem the subject generates non-random paths in the state-space representation of the problem. Solution paths generated by the problem solver tend to be goal-directed, and segments of solution paths also form portions of goal-directed paths.

Hypothesis 2. Given a decomposition of the state-space of a problem into subproblems satisfying the conditions for the reduction procedure described in Sec. IV.A, then (a) subproblem solution paths tend to be subgoal-directed, and (b) when subgoals are attained, the paths exit from the respective subproblems.

Hypothesis 3. Identifiable stages occur during problem-solving corresponding to the solutions of subproblems. That is, paths occur which do not constitute solutions (or else do not constitute the most direct solutions) to the problem, but which do constitute solutions to

all of the isomorphic subproblems (of a particular structure) entered by the path.

It may be that the validity of Hypotheses 2 and 3 depends on the particular way that the state-space of the problem is decomposed into subproblems, since such a decomposition is not unique.

Hypothesis 4. When two subproblems of a problem have isomorphic state-spaces, the problem solver's respective paths through these subproblems tend to be congruent.

Hypothesis 5. Given a symmetry group G of automorphisms of the state-space of a problem, there tend to occur successive paths congruent modulo G in the state-space. Such occurrences often culminate in the solving of the problem.

Elaborating on Hypothesis 5, it would be extremely interesting if given a symmetry group G for a problem, one could demonstrate stages corresponding to the acquisition of subgroups of G. Hypothesis 5 (symmetry acquisition) is suggestive of the "insight" phenomenon which changes the gestalt of the problem solver (Allport, 1955; Wertheimer, 1945).

These hypotheses are not to be regarded as a definitive list, but rather as preliminary and indicative of the kind of analysis of the effects of problem structure that is possible. While some of the hypotheses may seem intuitively obvious or necessary, it is not difficult to construct mechanical problem-solving mechanisms which violate any or all of them. Thus, if valid, they represent fairly general constraints on the properties which artificial intelligence models must

display in order to simulate humen problem-solving adequately.

The hypotheses focus on paths within the state-space, rather than on the operators that generate these paths. Different formal rules of procedure may sometimes lead to the same transitions or paths, just as different descriptions of a mathematical function may nevertheless define the same mapping. Thus to the extent that one seeks to describe behavior using operators which have precise domains of states and act on these within the state-space, the present formalism will be satisfactory. Hypothesis 3, for example, allows the interpretation of solution paths within isomorphism classes of subproblems as the application of a single operator which maps initial states of such subproblems into goal states.

However different descriptions of the same operator can of course imply different problem-solving strategies, just as different embodiments of a problem state-space can elicit different strategies. We are not seeking at this point to study the particular rules which subjects employ, but rather the structural features of the behavior they exhibit. This is the sense in which the present paper is concerned with problem-solving structures rather than with strategies.

In order to investigate empirically hypotheses such as the above, it seems natural to begin with a problem whose state-space possesses somewhat more symmetry than the problem environment presents perceptually, and which displays a rich subproblem structure. The Tower of Hanoi problem was selected for empirical investigation for these reasons. Classes of subproblems exist which are isomorphic to each other, and the state-space of each subproblem as well as that of the main problem possesses symmetry. In Section V the above hypotheses are interpreted

in terms specific to the Tower of Hanoi problem.

Let us remark again that we are regarding the problem as distinct from and exterior to the problem solver. In principle the state of the problem (more generally, the state of the environment) may be regarded as observable, the full state-space of the problem definable, and its structure ascertainable. On the other hand the state of the problem solver, and the mechanisms for change of that state, are to be inferred from the trace of the problem solver's behavior in the state-space of the problem. Thus we agree with the postulate of a feedback loop between the problem and the problem solver, rather than choosing to incorporate the two systems into a single information processing system (cf. Sec. II.B).

V. The Tower of Hanoi Problem

In this section we seek to make the foregoing ideas concrete by describing a problem that has been used for empirical investigation (Luger, 1973). The Tower of Hanoi is a well-known problem that has been extensively discussed (Dudeney, 1907; The Mathematics Teacher, 1951; Gardner, 1959). Its state-space is depicted by Nilsson (1971), and it has already been posed as a problem eminently suitable for mechanical problem-solving (Hormann and Shaffer, undated).

In the Tower of Hanoi problem as we study it, four concentric rings (labelled 1,2,3,4 respectively) are placed in order of size, the largest at the bottom, on the first of three pegs (labelled A,B,C); the apparatus is pictured in Figure 9. The object of the problem is to transfer all of the rings from peg A to peg C in a minimum number of moves. Only one ring may be moved at a time, and no larger ring may be placed above a smaller one on any peg.

Figure 10 is the complete state-space representation of the 4-ring Tower of Hanoi problem. Each circle stands for a possible position or state of the Tower of Hanoi board. The four letters labelling a state refer to the respective pegs on which the four rings are located. For example, state "CCBC" means that ring 1 (the smallest), ring 2 (the second smallest) and ring 4 (the largest) are in their proper order on peg C. Ring 3 (the next to the largest) is on peg B. Only states adjacent in the diagram are "connected" by the legal moves of the game; that is, a permitted move by the problem solver always effects a transition between states represented by neighboring circles in Figure 10. The solution path containing the minimum number of moves consists of

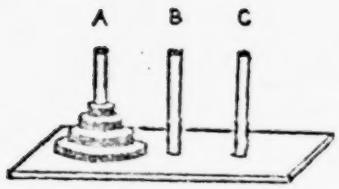

FIGURE 9. The Tower of Hanoi board in its initial state.

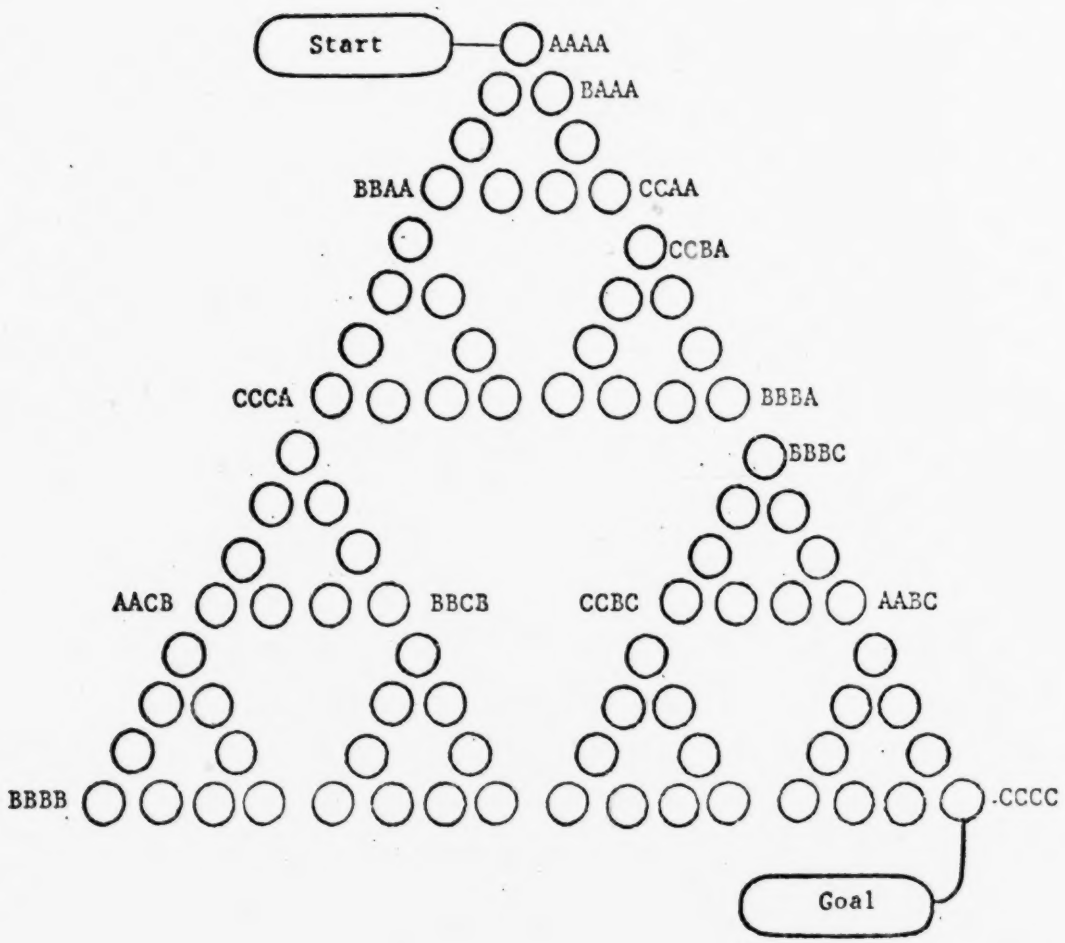

FIGURE 10. State-space representation of the 4-ring Tower of Hanoi problem.

The four letters labelling a state refer to the respective pegs on which the the four rings are located. Legal moves effect transitions between adjacent states.

the fifteen steps from state AAAA to state CCCC down the right-hand side of the state-space diagram.

A. Subproblem and Symmetry Structures

The Tower of Hanoi problem has a natural decomposition into nested subproblems, as pictured in Figure 11. In order to solve the 4-ring problem, it is necessary at some point to move the largest ring from its original position on peg A to peg C, but before this can be done the three smaller rings must be assembled in their proper order on peg B. The problem of moving the three rings from one peg to another may be termed a 3-ring subproblem, and constitutes a subset of the state-space representation.

As can be noted from Figure 11, the state-space for 4-ring Tower of Hanoi contains three isomorphic 3-ring subspaces, for which the physical problem-solving situations are different by reason of the position of ring 4 (the largest ring). Each subspace becomes a subproblem when one of its entry states is designated as the initial state, and its exit states are designated as goal states. The various 3-ring subproblems in turn differ from each other in that the rings are moved between different pegs, as well as with respect to the position of the largest ring.

Similarly each 3-ring subspace contains three isomorphic 2-ring subspaces, for a total of nine in the entire state-space. The various 2-ring subproblems differ with respect to the positions of rings 3 and 4, and with respect to the pegs between which the rings are to be moved.

Finally each 2-ring subspace may be further decomposed into three 1-ring subspaces, comprising only three states apiece.

Each n-ring subproblem, as well as the main problem, admits of a

FIGURE 11. Subproblem decomposition of the Tower of Hanoi state-space.

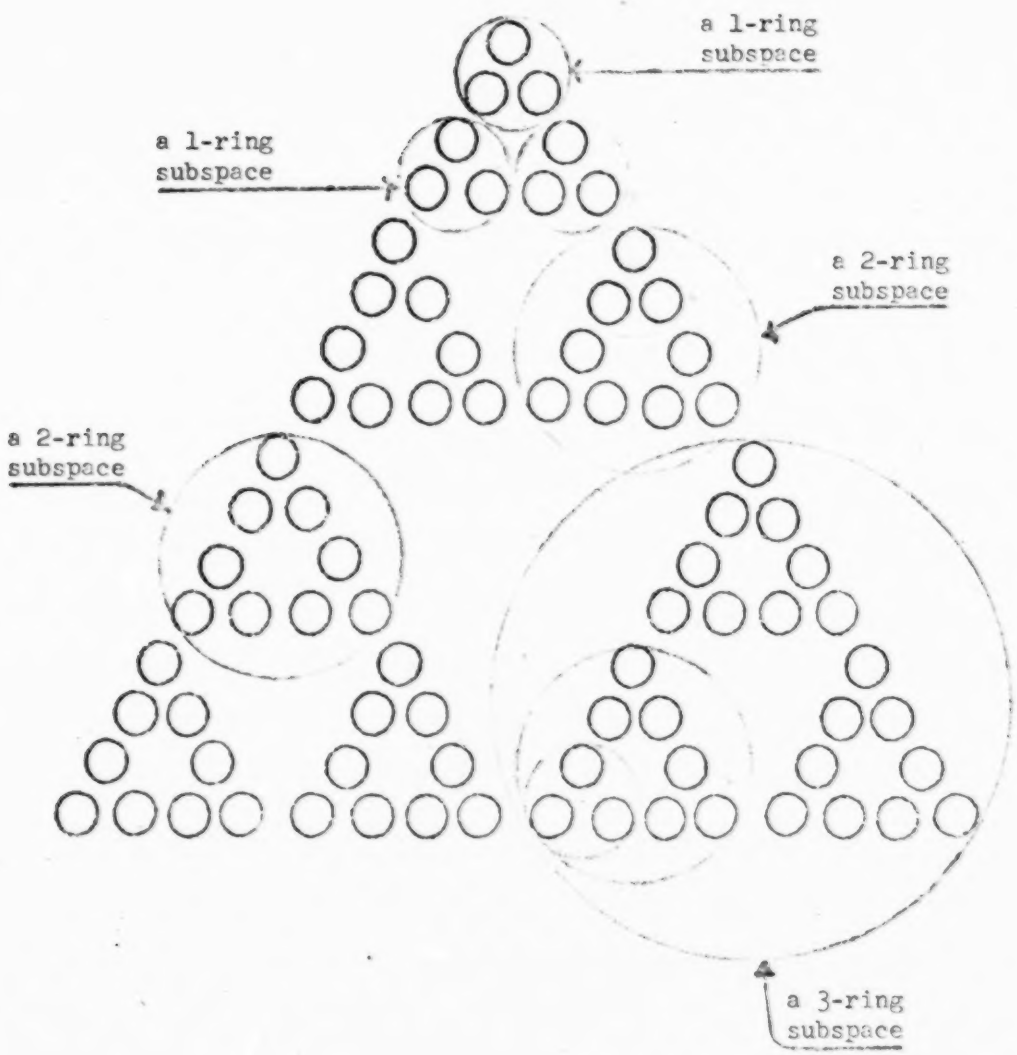

symmetry automorphism. The automorphism maps a goal state of the n-ring problem into the conjugate goal state which corresponds to transferring the n rings to the other open peg. Were the three pegs of the Tower of Hanoi board to be arranged at the corners of an equilateral triangle, the symmetry automorphism would represent the geometric operation of reflection about an altitude of the triangle.

B. Discussion of Hypotheses

Next we seek to interpret the hypotheses proposed in Section IV.B for subjects solving the Tower of Hanoi problem. To do so it is necessary to establish criteria for "goal-directed" paths through the state-space of the problem or one of its n-ring subproblems. We shall say that a solution path through such a state-space is goal-directed if the same state is not entered twice, and at each step the distance from the goal state (exit state) is non-increasing. The distance between two states in the Tower of Hanoi state-space is the minimum number of steps actually necessary to reach one state from the other.

Figure 12a illustrates the six mutually non-congruent goal-directed paths through a Tower of Hanoi 2-ring subproblem; three examples of non-goal-directed solution paths are given in Figure 12b.

For the 3- and 4-ring problems, it may be desirable to weaken the above criterion for goal-directedness by using a coarser measure of distance. Thus the distance between one state in the n-ring state-space and another can be defined as the minimum number of (n-2)-ring subspaces it is necessary to enter in order to reach the second state from the first. A solution path through the n-ring state-space is

FIGURE 12a. Goal-directed paths through a 2-ring Tower of Hanoi subproblem. There are six mutually non-congruent goal-directed paths.

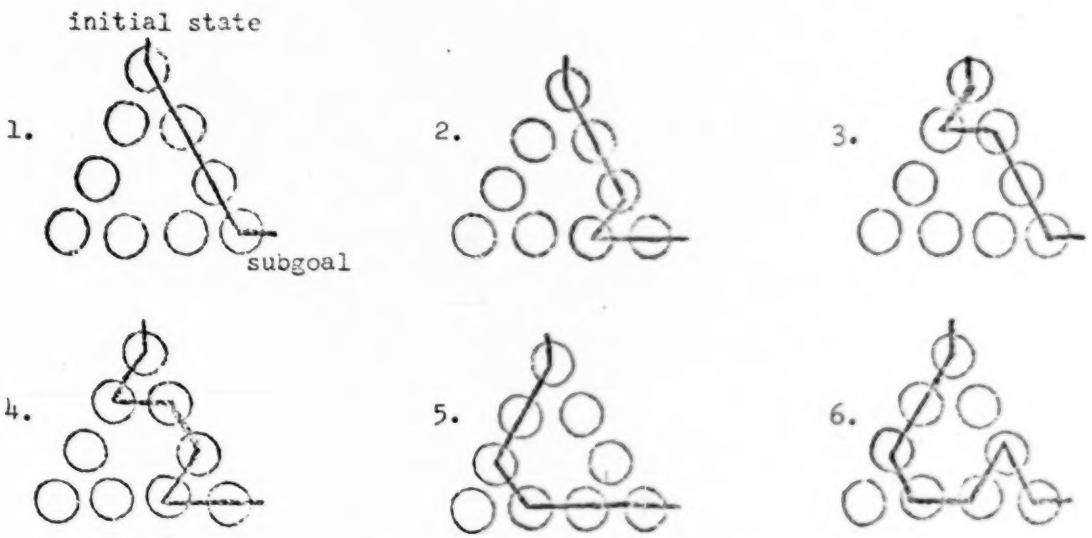

FIGURE 12b. Examples of non-goal-directed solution paths.

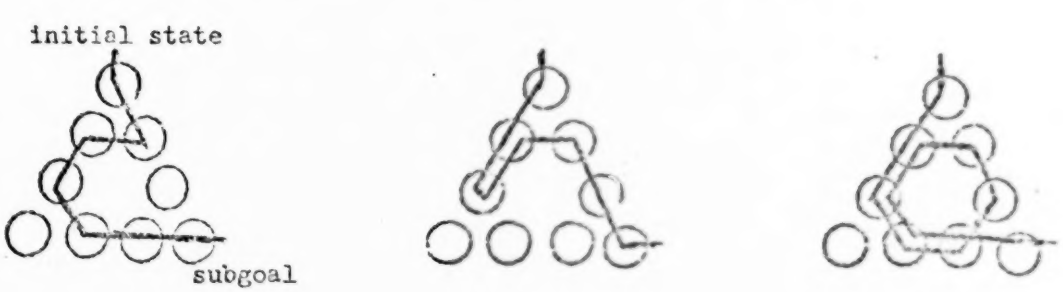

goal-directed (in the weaker sense) if the path does not enter the same (n-2)-ring subspace twice, and if the distance from the goal state is non-increasing. In short, "weak goal-directedness" corresponds to reducing the 3-ring state-space modulo its decomposition into 1-ring subproblems, and the 4-ring state-space modulo its decomposition into 2-ring subproblems, prior to determining the goal-directedness of a path. Thus it is possible for a path through the 3-ring subproblem to be subgoal-directed while a segment of the same path passing through a 2-ring subproblem is not. A path which is subgoal-directed within every subproblem traversed, as well as being goal-directed may be termed strongly goal-directed. The first diagram in Figure 13 illustrates a strongly goal-directed path; the path in the second diagram is not.

A randomly generated path is less likely to be goal-directed than might at first be supposed, even when one forbids the immediate retraction of a step. For the 2-ring subproblem such random paths are non-goal-directed 11 times in 32; the respective frequencies of occurrence of random paths congruent to the solution paths in Figure 12a are: $\frac{1}{4}, \frac{1}{8}, \frac{1}{8}, \frac{1}{16}, \frac{1}{16}, \frac{1}{32}$. Randomly generated paths are non-goal-directed more often in the case of the 3- and 4-ring problems.

Having distinguished "goal-directed" paths, Hypothesis 1 and Hypothesis 2a proposed in Section IV.B can be readily interpreted for the Tower of Hanoi problem.

Figure 14 illustrates the distinction in Hypothesis 2b between paths which do and do not exit from their respective subproblems when subgoal states are attained.

FIGURE 13. Goal-directed paths through a 3-ring Tower of Hanoi subproblem.

Note that the slight double-back in the second diagram does not disqualify the path from being goal-directed, since it occurs entirely inside a 1-ring subspace.

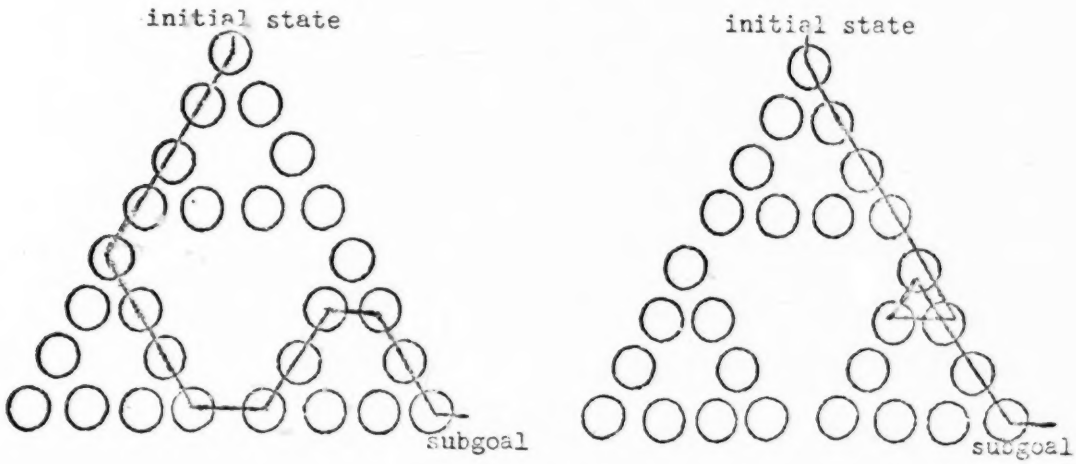

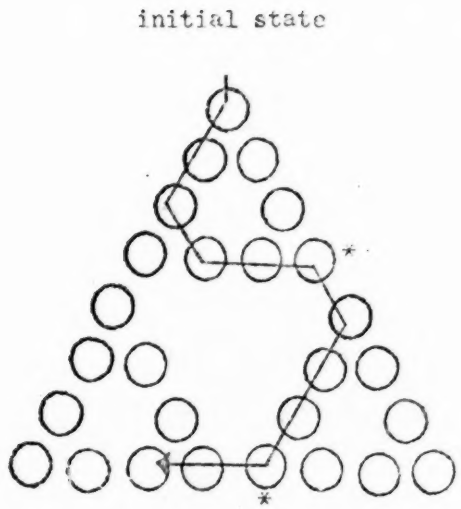

FIGURE 14a. A path illustrating Hypothesis 2b.

The path exits (*) from each 2-ring subproblem when a subgoal is attained.

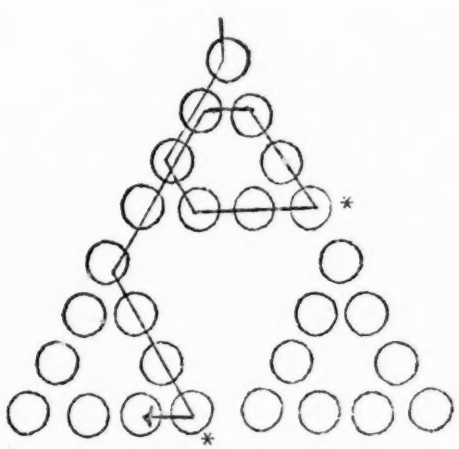

FIGURE 14b. A path violating Hypothesis 2b.

The path fails twice to exit (*) from a 2-ring subproblem when a subgoal state is attained.

Figure 15 depicts a state in problem-solving, as postulated in Hypothesis 3, where the 2-ring subproblem is consistently solved in the minimum number of steps, but the 3-ring subproblem is not. The state-space has been effectively reduced *modulo* its 2-ring subproblem decomposition.

Figure 16 gives several examples of congruent paths through isomorphic subproblems, as in Hypothesis 4.

Finally Figure 17 illustrates an instance of Hypothesis 5 - two successive paths through the Tower of Hanoi state-space congruent *modulo* the symmetry automorphism.

C. Subjects' Problem-Solving Behavior

In Figure 18 are pictured the actual paths through the state-space generated by two adult subjects solving the Tower of Hanoi problem. Each subject was shown the Tower of Hanoi board, and given verbal instructions. His moves on the board, as well as his comments and conversation with the investigator, were recorded on tape.

The behaviors of Subject A conform perfectly to all five of the proposed hypotheses. The paths are both goal- and subgoal-directed, and exit from a subproblem whenever a subgoal state is attained (Hypotheses 1 and 2). The first two trials contain five instances of solution of the 2-ring subproblem in the minimum number of steps, and two instances of solution in more than the minimum number, while the 3-ring subproblem has not yet been solved by the shortest path (Hypothesis 3). Trial 1 beautifully illustrates congruent paths through two isomorphic 3-ring subproblems (Hypothesis 4). Finally, trial 2 is interrupted (with the comment, "Could I try again? ... This is annoying

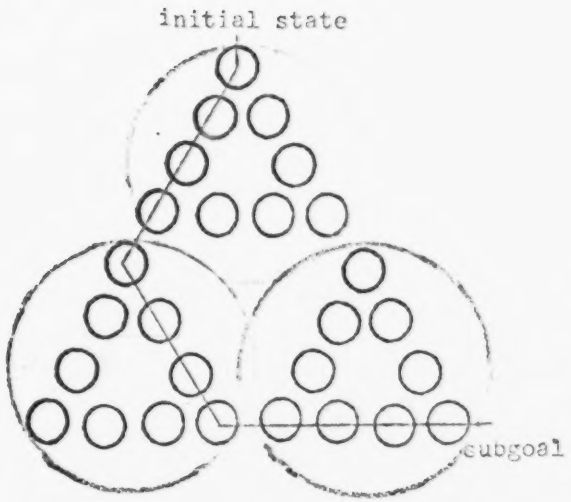

FIGURE 15. A stage in problem-solving.

The 2-ring subproblem is consistently solved in the minimum number of steps, while the 3-ring subproblem is not. The state-space has been effectively reduced modulo its 2-ring subproblem decomposition.

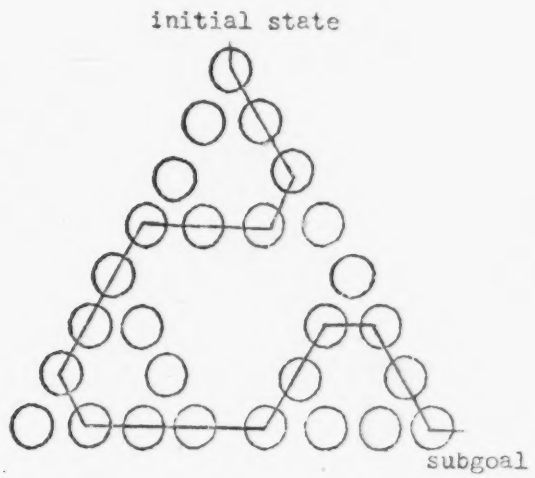

FIGURE 16. Congruent paths through isomorphic subproblems.

All three paths through the 2-ring subproblems in Figure 16 are congruent to each other.

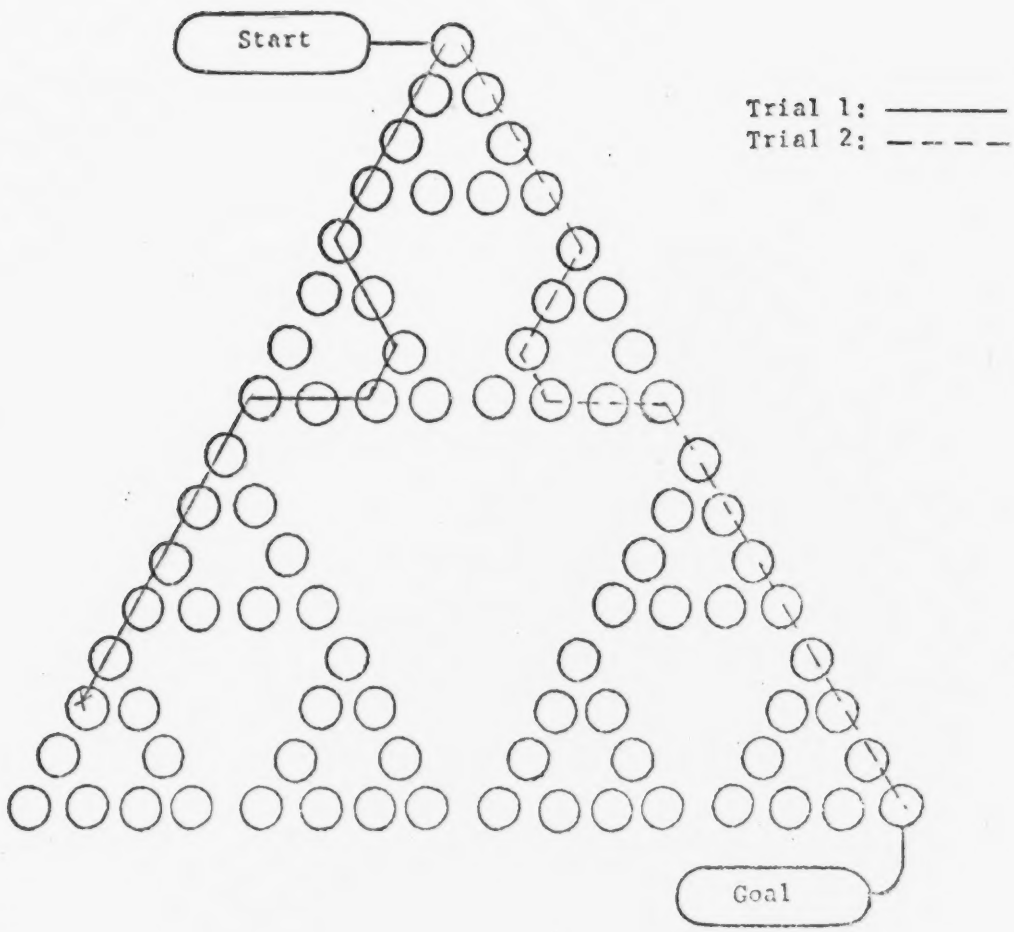

FIGURE 17. Two successive paths congruent modulo the symmetry automorphism.

FIGURE 18. Tower of Hanoi: Subjects' paths through state-space.

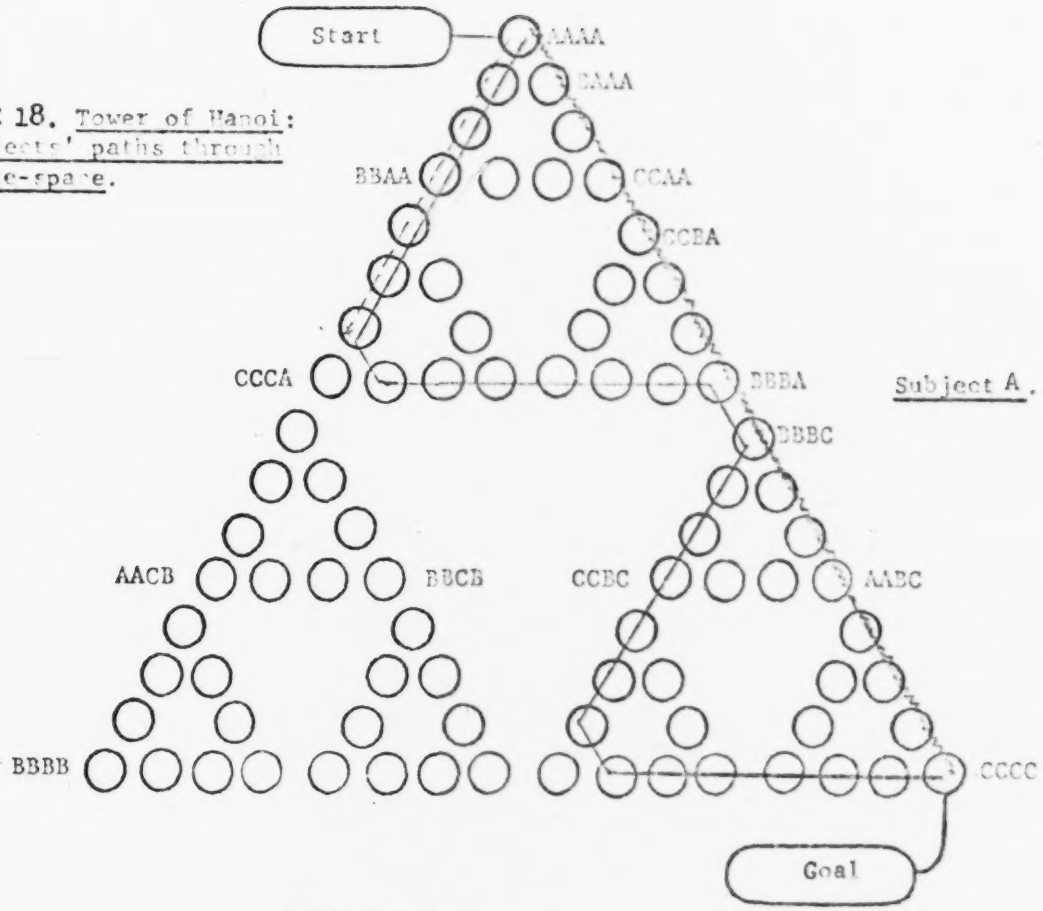

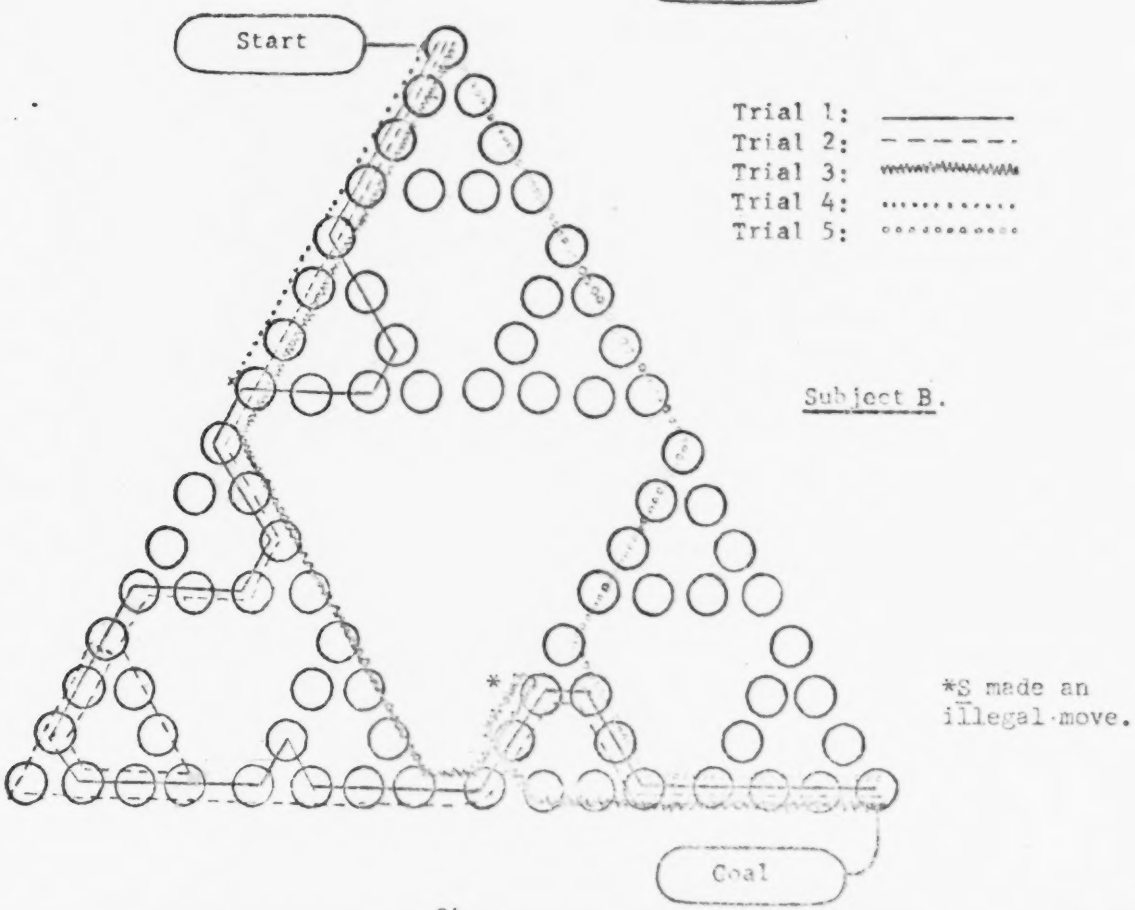

... "), and trial 3 (the shortest solution path) follows as the image of trial 2 under the symmetry automorphism - peg B ↔ peg C (Hypothesis 5).

Subject B displays more complicated problem-solving behavior. Both trials 2 and 3 contain segments which are not subgoal-directed (Hypothesis 2a), although they are goal-directed paths in the weaker sense (Hypothesis 1). During trial 3 the subject actually makes an illegal move. The second trial fails twice to exit from the 2-ring subproblem at the lower left of the diagram, despite having achieved "subgoal" states of that subproblem (Hypothesis 2b). The immediately prior comment was "I've lost my way." Thus the second hypothesis is not fully satisfied, although it is far more closely obeyed than would be expected from randomly generated paths.

In trials 3-5, the subject solves the 3-ring subproblem by the shortest path four out of six times and the 2-ring subproblem twelve out of thirteen times, while the 4-ring problem has not yet been mastered (Hypothesis 3). Trials 1 and 2 display seven examples of congruent paths (type 5, Figure 12a) through isomorphic 2-ring subproblems, as well as nine instances of the shortest path (type 1) and one instance of a path (type 3) non-congruent to the others (Hypothesis 4).

Finally trial 4 is interrupted, and trial 5 follows as its image under the symmetry automorphism; the subject deviates from the shortest solution path only after completing the segment congruent to trial 4 (Hypothesis 5).

Thus we have seen in the case of the paths generated by two individual problem solvers, illustrations of and a considerable degree of adherence to the suggested hypotheses.

A paper to follow by one of the authors (GFL) tabulates the behaviors of forty-five adult subjects solving the Tower of Hanoi problem, testing more systematically the validity of the five hypotheses we have proposed.

VI. Conclusion

Several authors have recently sought to distinguish "strategy" from "structure" in problem-solving, and to investigate the relationship between them (Dienes and Jeeves, 1965, 1970; Branca and Kilpatrick, 1972). The present paper suggests one natural way to make this distinction. We let the structure of a problem refer to the formal properties of its state-space representation, such as the symmetry automorphisms which are present and the subproblem decompositions which are possible. We consider the subject's cognitive structures to include the conservation operations, symmetries, and subproblem decompositions that he perceives in the problem situation. These determine the states that he treats as distinct and those he treats as equivalent. They may change during the course of problem-solving, leading to effective reduction of the state-space. The subject's behavior can be faithfully mapped as long as the state-space representation that is utilized by the researcher is sufficiently detailed, in that it does not treat states as equivalent which the subject treats as distinct.

We let the term strategy refer to particular rules or procedures for taking steps within the state-space once the latter has been established. These are analogous to the search algorithms employed in mechanical problem-solving. Different individuals clearly use different strategies in solving the same problem, and the same individual often employs different strategies in solving different but isomorphic problems. The present paper does not examine strategies per se, but hypothesizes that even in the context of different strategies, certain patterns in behavior tend to occur as a consequence of

acquisition of elements of the problem structure by the subject. Artificial intelligence models will have to incorporate such structural effects in order to adequately simulate human problem-solving behavior.

Acknowledgements

The authors would like to thank John Carr of the Moore School of Electrical Engineering and Gail Zivin of the Graduate School of Education at the University of Pennsylvania for many interesting discussions and valuable suggestions.

References

Allport, F.H., Theories of Perception and the Concept of Structure. New York: Wiley, 1955.

Arbib, M.A., Algebraic Theory of Machines, Languages and Semigroups. New York: Academic Press, 1969.

Banerji, R.B., Theory of Problem Solving. New York: American Elsevier Publishing Co., 1969.

Bartlett, F.C., Thinking. New York: Basic Books, 1958.

Bourbaki, N. Elements of Mathematics - Series. Reading, Mass: Addison-Wesley.

Branca, N.A., and Kilpatrick, J. "The Consistency of Strategies in the Learning of Mathematical Structures," Journal for Research in Mathematics Education, Vol. 3, No. 3 (May 1972).

Bruner, J.S., Goodnow, J.J., and Austin, G.A., A Study of Thinking. New York: Wiley, 1956.

Carr, J.W., Unpublished Lecture Notes, Moore School of Electrical Engineering, University of Pennsylvania, Philadelphia, Pennsylvania.

Dienes, Z.P., and Jeeves, M.A., Thinking in Structures. London: Hutchinson Educational, 1965.

Dienes, Z.P., and Jeeves, M.A., The Effects of Structural Relations on Transfer. London: Hutchinson Educational, 1970.

Dudeney, E., The Canterbury Puzzles, 1907.

Eisenbud, L., and Wigner, E.P., Nuclear Structure. Princeton, N.J.: Princeton University Press, 1958.

Feynman, R.P., Lectures on Physics, Vol. III. Reading, Mass: Addison-Wesley, 1965.

Gamow, G., One, Two, Three ... Infinity. New York: Viking Press, 1947.

Gardner, M., Mathematical Puzzles & Diversions. New York: Simon and Schuster, 1959.

Gelernter, H., "A Note on Syntactic Symmetry a the Manipulation of Formal Systems by Machine," Information and Control, 2, 80-89 (1959).

Gelernter, H., "Realization of a Geometry Theorem Proving Machine," Proceedings of 1959 International Conference on Information Processing, 273-282. Paris: UNESCO, 1960.

Ginsburg, H., and Opper, S., Piaget's Theory of Intellectual Development: an Introduction, Englewood Cliffs, N.J.: Prentice-Hall, 1969.

References (continued)

Guzman, A., "Computer Recognition of 3-Dimensional Objects in a Visual Scene," Cambridge: Massachusetts Institute of Technology, MAC-TR-59, 1968.

Harris, Z. S., Structural Linguistics. Chicago: The University of Chicago Press, 1951.

Hormann, A.M., and Shaffer, S.S., "Problem Solving by Artificial Intelligence." Santa Monica California: Systems Development Corp., undated manuscript.

Johnson, E.S., "An Information Processing Model of One Kind of Problem Solving," Psychological Monographs, 78, No. 4 (1964).

Kimble, G.A., Hilgard & Marquis' Conditioning and Learning. New York: Appleton-Century-Crofts, Inc., 1961.

Lane, M., Introduction to Structuralism. New York: Basic Books, 1970.

Lévi-Strauss, C., Structural Anthropology. New York: Basic Books, 1963.

Lévi-Strauss, C., Elementary Structures of Kinship. Boston: Beacon, 1969.

Luger, G.F., Ph.D. Thesis, Graduate School of Education, University of Pennsylvania, 1973.

Mathematics Teacher, 44, p. 505 (1951).

Menzel, W., Theorie der Lernsysteme. Springer Verlag: Berlin, 1970.

Minsky, M., & Papert, S., Progress Report of the MIT Artificial Intelligence Lab. Boston: MIT Press, 1972.

Newell, A., Shaw, J.C., & Simon, H.A., "Report on a General Problem Solving Program." The RAND Corporation Paper P-1584 (Feb. 1959).

Newell, A., & Simon, H.A., Human Problem Solving. Englewood Cliffs, N.J.: Prentice-Hall, 1972.

Nilsson, N., Problem Solving Methods in Artificial Intelligence. New York: McGraw-Hill, 1971.

Piaget, J., Structuralism. New York: Basic Books, 1970.

Piaget, J., & Inhelder, B. The Psychology of the Child. New York: Basic Books, Inc., 1969.

Polya, G., How to Solve It. Princeton, N.J.: Princeton University Press, 1945.

References (continued)

Tolman, E.C., "Cognitive Maps in Rats and Men," Psychology Review, 55, 189-209 (1948).

Wertheimer, M., Productive Thinking. New York: Harper & Brothers, 1945.

Wigner, E.P., Group Theory. New York: Academic Press, 1959.

Wigner, E.P., "Symmetry and Conservation Laws," Proceedings of the National Academy of Sciences, 51, No. 5 (May 1964).

Winograd, T., "Procedures as a Representation for Data in a Computer Program for Understanding Natural Language," Cambridge: Massachusetts Institute of Technology, AI-TR-17, 1971.

APPENDIX

Behaviors of Subjects Solving the Tower of Hanoi Problem

by

George F. Luger

February 1973

TABLE 1. Goal-directed paths and path segments generated by subjects solving the 4-ring Tower of Hanoi problem (test of Hypothesis 1). See next page for totals and key to notation.

Subject	Trial Number								Totals				Fraction strongly goal-directed	Fraction all goal-directed
	1	2	3	4	5	6	7	8	*	+	-	sum		
1	*	+	+	*	*				3	2	0	5	.60	1.00
2	+	-	*	1					1	1	1	3	.33	.67
3	*	1							1	0	0	1	1.00	1.00
4	+	1							0	1	0	1	.00	1.00
5	*	1							1	0	0	1	1.00	1.00
6	*	*	1						2	0	0	2	1.00	1.00
7	*	*	1						2	0	0	2	1.00	1.00
8	*	1							1	0	0	1	1.00	1.00
9	*	1							1	0	0	1	1.00	1.00
10	-	1							0	0	1	1	.00	.00
11	+	*	*	1					2	1	0	3	.67	1.00
12	*	*	*	*	1				4	0	0	4	1.00	1.00
13	+	*	*	+	*	1			3	2	0	5	.60	1.00
14	*	*	*						3	0	0	3	1.00	1.00
15	+	*							1	1	0	2	.50	1.00
16	*	1							1	0	0	1	1.00	1.00
17	*	1							1	0	0	1	1.00	1.00
18	*	*	+	*	*				4	1	0	5	.80	1.00
19	*	*	1						2	0	0	2	1.00	1.00
20	-	*	-	1					1	0	2	3	.33	.33
21	*	1							1	0	0	1	1.00	1.00
22	*	+	*						2	1	0	3	.67	1.00
23	-	*	*	1					2	0	1	3	.67	.67
24	*	+	*	1					2	1	0	3	.67	1.00
25	-	*	*	+	*	+	*	1	4	2	1	7	.57	.86
26	*	-	*	1					2	0	1	3	.67	.67
27	*	*	*	1					3	0	0	3	1.00	1.00
28	*	*	1						2	0	0	2	1.00	1.00
29	+	*	*	+	*	*	*		5	2	0	7	.71	1.00
30	-	-	+	+	1				0	2	2	4	.00	.50
31	*	*	-	1					2	0	0	2	1.00	1.00
32	*	1							1	0	0	1	1.00	1.00
33	*	*	*	1					3	0	0	3	1.00	1.00
34	-	*	+	*	1				2	1	1	4	.50	.75
35	*	*	1						2	0	0	2	1.00	1.00
36	*	+	*	*	*	*	1		5	1	0	6	.83	1.00
37	*	+	*	*	*	1			4	1	0	5	.80	1.00
38	*	*	*	1					3	0	0	3	1.00	1.00
39	*	*	*	1					3	0	0	3	1.00	1.00
40	*	*	*	+	*				4	1	0	5	.80	1.00
41	*	*	*						3	0	0	3	1.00	1.00
42	*	*	*	*	*	*	*	1	7	0	0	7	1.00	1.00
43	*	*	*	1					3	0	0	3	1.00	1.00
44	*	*	1						2	0	0	2	1.00	1.00
45	*	1							1	0	0	1	1.00	1.00

TABLE 1 (continued).

	Totals				Fraction strongly goal-directed	Fraction all goal-directed
	*	+	-	sum		
Total all subjects, all trials	102	21	10	133	.77	.93
Total all subjects, first trial only	33	6	6	45	.73	.87

Key

1 = most direct solution path (not included in sum total)
* = strongly goal-directed path (goal-directed within each subproblem), excepting the most direct solution path
+ = goal-directed path, but not strongly goal-directed
- = non-goal-directed path

Among paths generated by the random selection of successor states, in which the retraction of a move is forbidden, less than 2/3 are goal-directed and less than 10% strongly goal-directed.

TABLE 2. Paths of subjects through 2-ring subproblems of the Tower of Hanoi.

Key

1, ..., 6 = congruence class of goal-directed path
X = non-goal-directed path (Hypothesis 2a)
- = failure to exit from subgoal state of subproblem (Hypothesis 2b)
C = significant congruence among goal-directed paths through isomorphic subproblems
NC = significant non-congruence among goal-directed paths through isomorphic subproblems (Hypothesis 4)

Subject	Sequence of 2-ring subproblems	con-gruence	1	2	3	4	5	6	X	total	-
1	1 5 5 5 3 5 1 1 1 5 1-5-1-1 1 5 1 1 1 1 1 X 1 1 1 1 1 1 1 1	C	21	0	1	0	7	0	1	30	3
2	1 1-1 1 1 1 1 1 1 1 1 1 1 1		14	0	0	0	0	0	0	14	1
3	5 1 1 1 1 1 1 1 1		8	0	0	0	1	0	0	9	0
4	1 X-1 1 1 1 1 1 1		8	0	0	0	0	0	1	9	1
5	1 1 1 1 1 1 1 1 1		9	0	0	0	0	0	0	9	0
6	1 1 1 1 X 1 1 1 1 1		9	0	0	0	0	0	1	10	0
7	1 1 1 1 1 5 1 1 1 1 1 1 1		12	0	0	0	1	0	0	13	0
8	1 2 1 1 1 1 1 1 1		8	1	0	0	0	0	0	9	0
9	1 1 1 1 1 1		6	0	0	0	0	0	0	6	0
10	5 X 4 2 1 1 1 1 1-1 1 1 1 1 1 1 1	NC	13	1	0	1	1	0	1	17	1
11	1 1-1-1 1 1 1 1 5 1 1 1 1 1 1 1 1 1		15	0	0	0	1	0	0	16	2
12	1 1 1 1 1 1 1 1 1 1		10	0	0	0	0	0	0	10	0
13	1 X 1 1 1 1 3 1 5 1 1 X 1 1 5 1 1 1 1 1		15	0	1	0	2	0	2	20	0
14	1 1 5 1 1 3 5 1 1 1 5 1	C	8	0	1	0	3	0	0	12	0
15	1 1 6 5 1 1 X 1 1 1 3 1	NC	8	0	1	0	1	1	1	12	0
16	1 2 1 1 3 1 1 1 1 1		8	1	1	0	0	0	0	10	0
17	1 5 3 1 1 1 1 1		6	0	1	0	1	0	0	8	0
18	1 3 1 1 1 1 1 3 3 1 1 X 1 3 1 1 1 3 1 1	C	14	0	5	0	0	0	1	20	0

TABLE 2 (continued).

Subject	Sequence of 2-ring subproblems	con-gruence	1	2	3	4	5	6	X	total	-
19	1 5 1 1 5 1 1 1 1 1 1 1		10	0	0	0	2	0	0	12	0
20	3 2 X 1 5 1 1 1 1 1 1 1 5 1 1 1 1 1 1 1 1 1	NC	18	1	1	0	2	0	1	23	0
21	6 1 1 1-1 1-1 3 1 1 1 1 1		11	0	1	0	0	1	0	13	2
22	1 1 1 1 1 1 X 1 1 3 1		9	0	1	0	0	0	1	11	0
23	1 1 1 5 1 1 3 1 5 1 1 5 1 5 1 1 1 1 1	C	14	0	1	0	4	0	0	19	0
24	1 1 1-1 1 1 X 1 1 1 1 1 1 1 1 1 1 1 1		18	0	0	0	0	0	1	19	1
25	5 X 3 X 1 1 1 1 5 1 1 1 1 5 1 1 X 1 1 X-X-1 1 5 1 1 1 1 1 1 1	C	21	0	1	0	4	0	5	31	2
26	1 1 1 X 1 1 1 1 1		8	0	0	0	0	0	1	9	0
27	1 1 1 1 3 1 1 1 1 5 1 1 1 1 1 1 1		15	0	1	0	1	0	0	17	0
28	1 3 5 1 1 1 1 1 1 1		8	0	1	0	1	0	0	10	0
29	1 1 1-1 1 1 2 5 1 1 1 1 1 1 1-X-1 1 1 1 1 1 1 1 3 1	NC	22	1	1	0	1	0	1	26	3
30	1 1-X 1 1 1 3 5 5-3 1-1-X 5 1 2 1 1 1 X-1 1 1 1 5 1 X 1 1 1 1 1	NC	21	1	2	0	4	0	4	32	5
31	1 3 1 3 1 1 1 1 1 5-1 1 1 1 1		12	0	2	0	1	0	0	15	1
32	1 1 1 1 1 1		6	0	0	0	0	0	0	6	0
33	1 3 2 1 1 5 1 1 1 1 1 1 1	NC	10	1	1	0	1	0	0	13	0
34	1 3 X-X 1 1 1 1 1 1 1 1 1 1 1 3 1 X 1 1 1 1 1 1 1 1 1 1 1		24	0	2	0	0	0	3	29	1
35	1 1 1 1 1 1 1 1 1 1 1		11	0	0	0	0	0	0	11	0
36	1 1 1 3 1 1 X 1 1 1 3 1-1 1 5 1 1 1 1 1 1		17	0	2	0	1	0	1	21	1
37	1 1 1 5 1 1 5 X 1 1 5 1 1 1 1 1 1 1 1	C	15	0	0	0	3	0	1	19	0
38	1 5 1 1 1 1-1 1 1 1 1		10	0	0	0	1	0	0	11	1
39	1 1 5 1 1 1 1 1 1 1 1 1 1 1 1		14	0	0	0	1	0	0	15	0

TABLE 2 (continued).

Subject	Sequence of 2-ring subproblems	con-gruence	1	2	3	4	5	6	X	total	-
40	1 1-1 1 1 1 1 1 1 4 1 1-1 1 1 3 1		15	0	1	1	0	0	0	17	2
41	1 1 5 1 1 1 5 1 1 1 2 1		9	1	0	0	2	0	0	12	0
42	1 1 1 1 1 1 1 1 1 1 1 5 6 1 1 1 1 1 1 5 1 1 1 1 1 1 1 1 1 1 1		27	0	0	0	2	1	0	30	0
43	1 1 5 5 5 5 1 1 1 1 1	C	7	0	0	0	4	0	0	11	0
44	1 1 1 1-1 1 1 1 1 1 1 1 1		13	0	0	0	0	0	0	13	1
45	1 1 1 1 1 1		6	0	0	0	0	0	0	6	0

				1	2	3	4	5	6	X	total	-
Totals	45 Subjects	7 C	6 NC	563	8	29	2	53	3	27	685	28

TABLE 3. Paths of subjects through 3-ring subproblems of the Tower of Hanoi.

Key

A, B, C, ... = congruence class of goal-directed path (specific to each subject)
X = non-goal-directed path (Hypothesis 2a)
- = failure to exit from subgoal state of subproblem (Hypothesis 2b)
C = significant congruence among goal-directed paths through isomorphic subproblems
NC = significant non-congruence among goal-directed paths through isomorphic subproblems (Hypothesis 4)

Subject	Sequence of 3-ring subproblems	con-gruence	1	A	B	C	D	E	F	X	total	-
1	A X B 1 X-1 B 1 1 X 1 1 C	NC	6	1	2	1				3	13	1
2	1-1 A A 1 1 1	C	5	2						0	7	1
3	X 1 1 1		3							1	4	0
4	X 1 1 1		3							1	4	0
5	A 1 1 1		3	1						0	4	0
6	1 1 A 1 1		4	1						0	5	0
7	1 1 A 1 1 1		5	1						0	6	0
8	A B 1 1	NC	2	1	1					0	4	0
9	A 1 1		2	1						0	3	0
10	X X 1 1 1		3							2	5	0
11	1-A 1 1 B 1 1 1	NC	6	1	1					0	8	1
12	A 1 1 A 1 1	C	4	2						0	6	0
13	X 1 1 A B X 1 B 1 1		5	1	2					2	10	0
14	1 A B A 1 A	C	2	3	1					0	6	0
15	1 X X 1 A		2	1						2	5	0
16	A 1 B 1 1	NC	3	1	1					0	5	0
17	A B 1 1	NC	2	1	1					0	4	0
18	A 1 B C D C 1 E 1	NC	3	1	1	2	1	1		0	9	0
19	A A 1 1 1	C	3	2						0	5	0
20	A X B C 1 B 1 D 1 1	NC	4	1	2	1	1			1	10	0
21	A 1 B C 1 1	NC	3	1	1	1				0	6	0

TABLE 3 (continued).

Subject	Sequence of 3-ring subproblems	con-gruence	1	A	B	C	D	E	F	X	total	-
22	A 1 X 1 B	NC	2	1	1					1	5	0
23	1 X A B C B 1 1	NC	3	1	2	1				1	8	0
24	1 A X 1 B B 1 1		4	1	2					1	8	0
25	X 1 1 A 1 1 A X X 1 A 1 1 1	C	8	3						3	14	0
26	1 X 1 1 1		4							1	5	0
27	1 A 1 B 1 1 1 1	NC	6	1	1					0	8	0
28	A B 1 1 1	NC	3	1	1					0	5	0
29	1 A 1 B C 1 1 X D 1 1 1 E	NC	7	1	1	1	1	1		1	13	0
30	1-X A B-C-X D X 1 1 E X 1 1	NC	5	1	1	1	1	1		4	14	3
31	A B 1 X 1 1	NC	3	1	1					1	6	0
32	1 1 1		3							0	3	0
33	A B C 1 1 1 1	NC	4	1	1	1				0	7	0
34	A X X 1 B B 1 C D B B 1 1	C	4	1	4	1	1			2	13	0
35	1 1 1 1 1 1		6							0	6	0
36	1 A B C D C E F 1 1	NC	3	1	1	2	1	1	1	0	10	0
37	A B C D C 1 1 1 1 1	NC	5	1	1	2	1			0	10	0
38	A 1 B 1 1	NC	3	1	1					0	5	0
39	1 A 1 1 1 1 1 1		7	1						0	8	0
40	1 1 A 1 B 1 1 1 C	NC	6	1	1	1				0	9	0
41	1 A 1 A 1 B		3	2	1					0	6	0
42	1 A 1 B B C 1 1 D 1 A A 1 1	NC	7	3	2	1	1			0	14	0
43	1 1 X A 1 1		4	1						1	6	0
44	1 A 1 1 1 1		5	1						0	6	0
45	1 1 1		3							0	3	0
Totals	45 Subjects 6 C 21 NC									28	321	6

TABLE 4. Subjects solving the Tower of Hanoi problem--stages corresponding to acquisition of the most direct solutions to 2- and 3-ring subproblems respectively (test of Hypothesis 3).

Key

1 = most direct solution to subproblem
0 = not the most direct solution

First row: sequence of 2-ring subproblems entered
Second row: sequence of 3-ring subproblems entered
Third row: sequence of trials on the 4-ring problem

! = acquisition of the most direct solution. This is defined to occur at the earliest point beyond which more than half of any sequence of subproblems (of the appropriate level) are solved in the most direct way.

* = point between successive paths through the 4-ring state-space congruent by virtue of a symmetry automorphism (Hypothesis 5)

Subject

1 2-ring: 1 0 0 0 0 0 ! 1 1 1 0 1 0 1 1 1 0 1 ! 1 1 1 1 1 1 0 1 * 1 1 1 1 1 !
3-ring: 0 0 0 1 0 1 0 ! 1 1 1 0 1 0 !
4-ring: 0 0 0 0 0 !

2 2-ring: ! 1 1 1 1 1 1 ! 1 1 ! 1 1 1 1 *
3-ring: 1 0 0 ! 1 ! 1 1
4-ring: 0 0 0 ! 1

3 2-ring: 0 ! 1 1 ! 1 1 ! 1 1 1 1
3-ring: 0 ! 1 ! 1 1
4-ring: 0 ! 1

4 2-ring: 1 0 ! 1 ! 1 1 ! 1 1 1 1
3-ring: 0 ! 1 ! 1 1
4-ring: 0 ! 1

5 2-ring: ! 1 1 1 ! 1 ! 1 1 1 1
3-ring: 0 ! 1 ! 1 1
4-ring: 0 ! 1

6 2-ring: ! ! 1 1 * 1 1 0 1 ! 1 1 1 1
3-ring: ! 1 1 0 ! 1 1
4-ring: 0 0 ! 1

7 2-ring: ! ! 1 1 1 1 1 0 1 1 1 ! 1 1 1 1 *
3-ring: ! 1 1 0 1 ! 1 1
4-ring: 0 0 ! 1

TABLE 4 (continued).

Subject

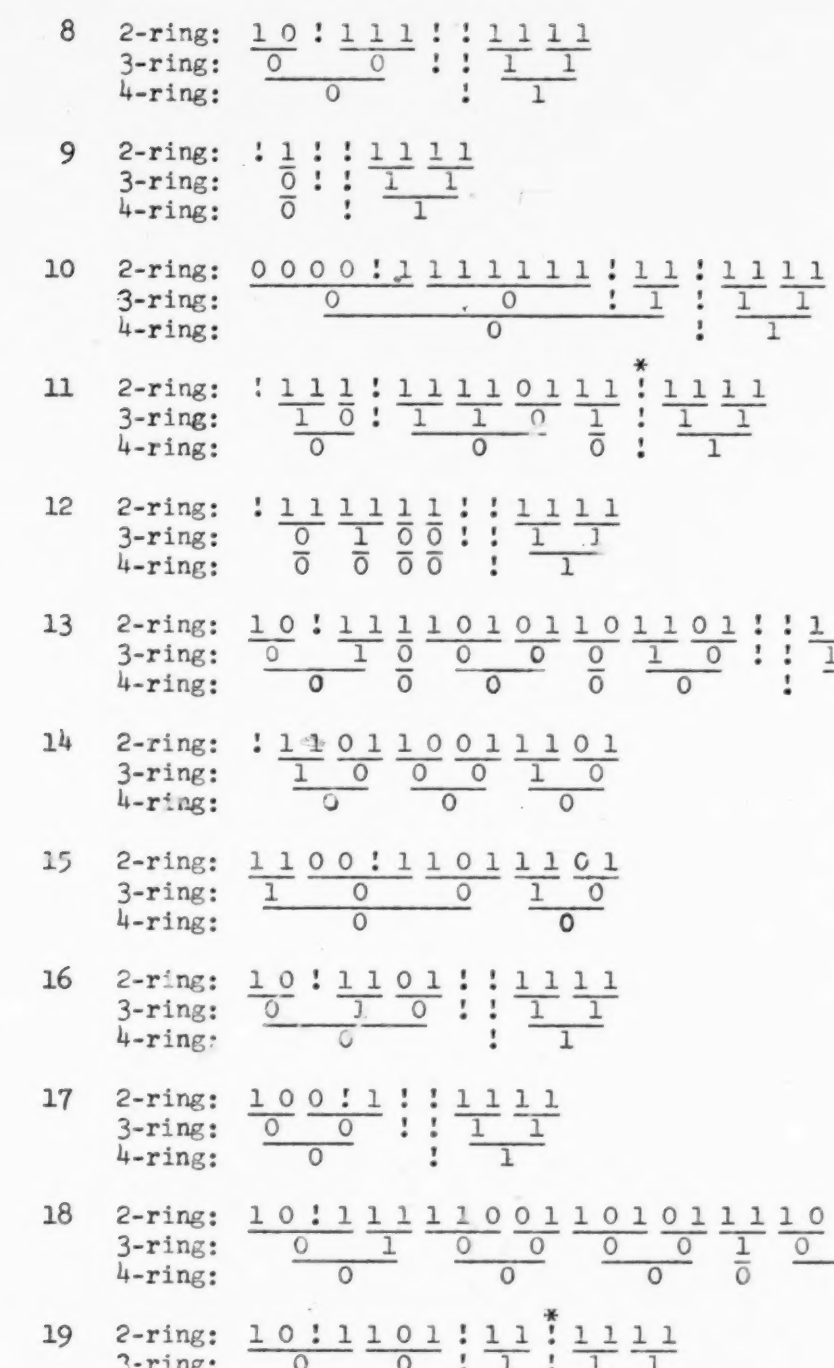

TABLE 4 (continued).

Subject

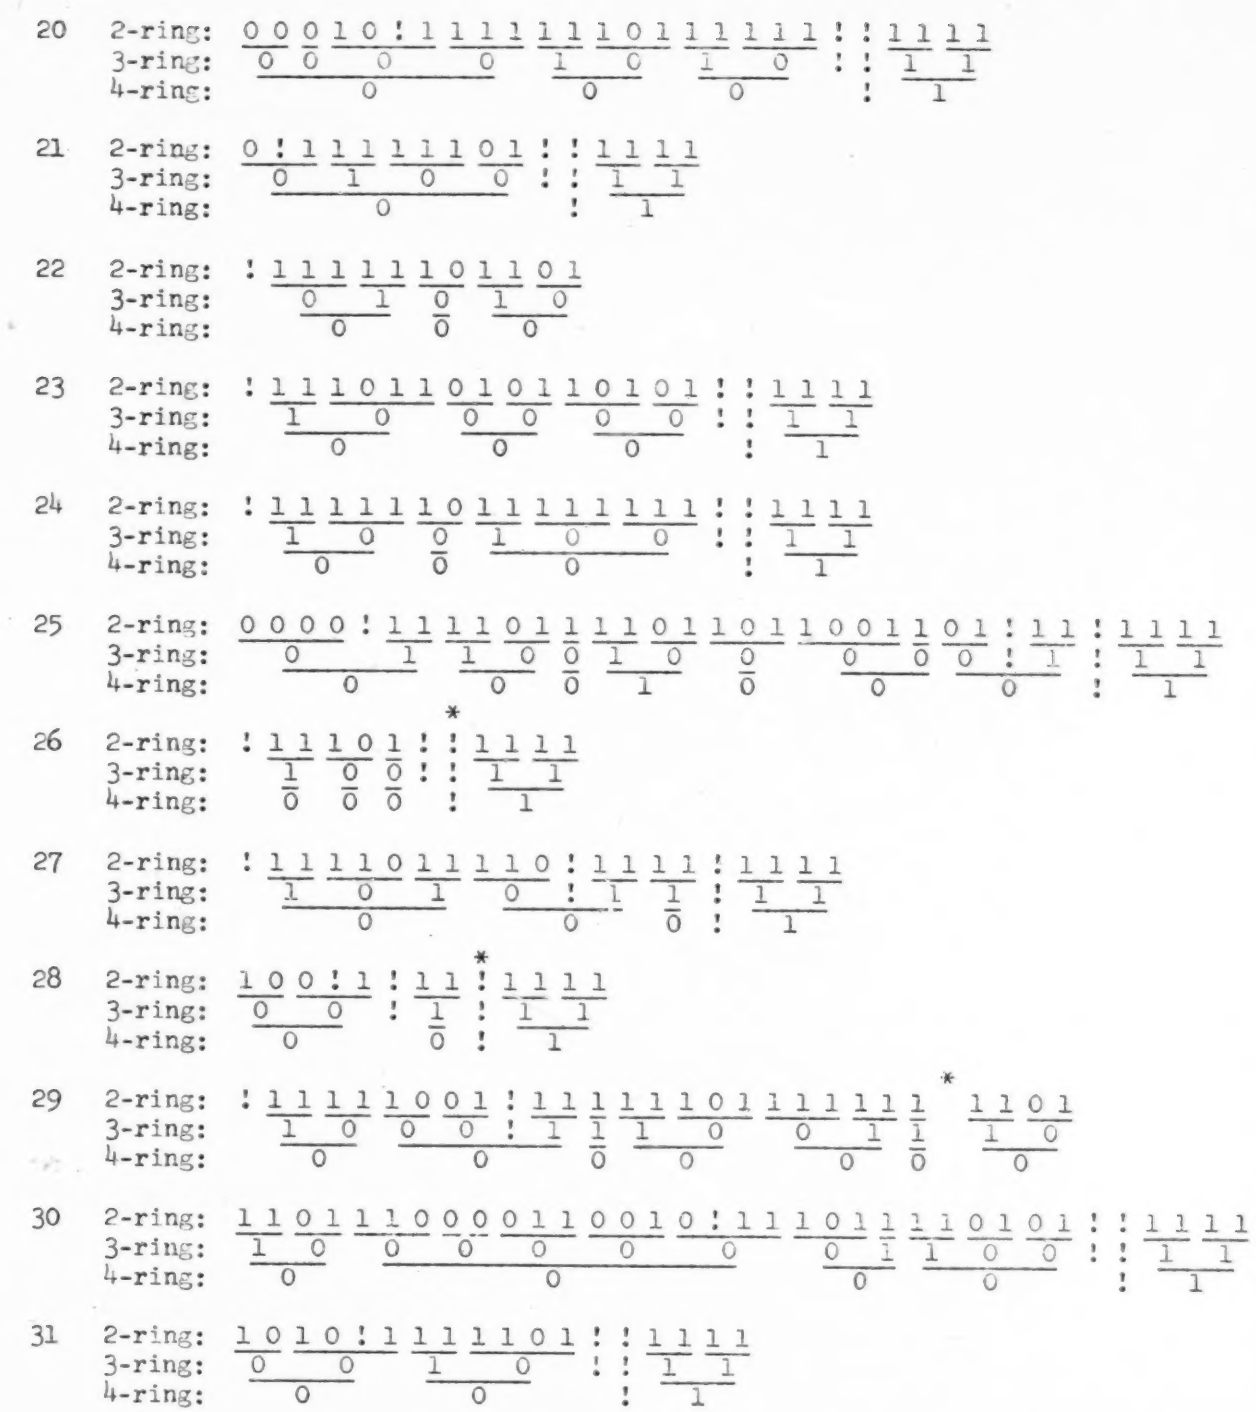

TABLE 4 (continued).

Subject

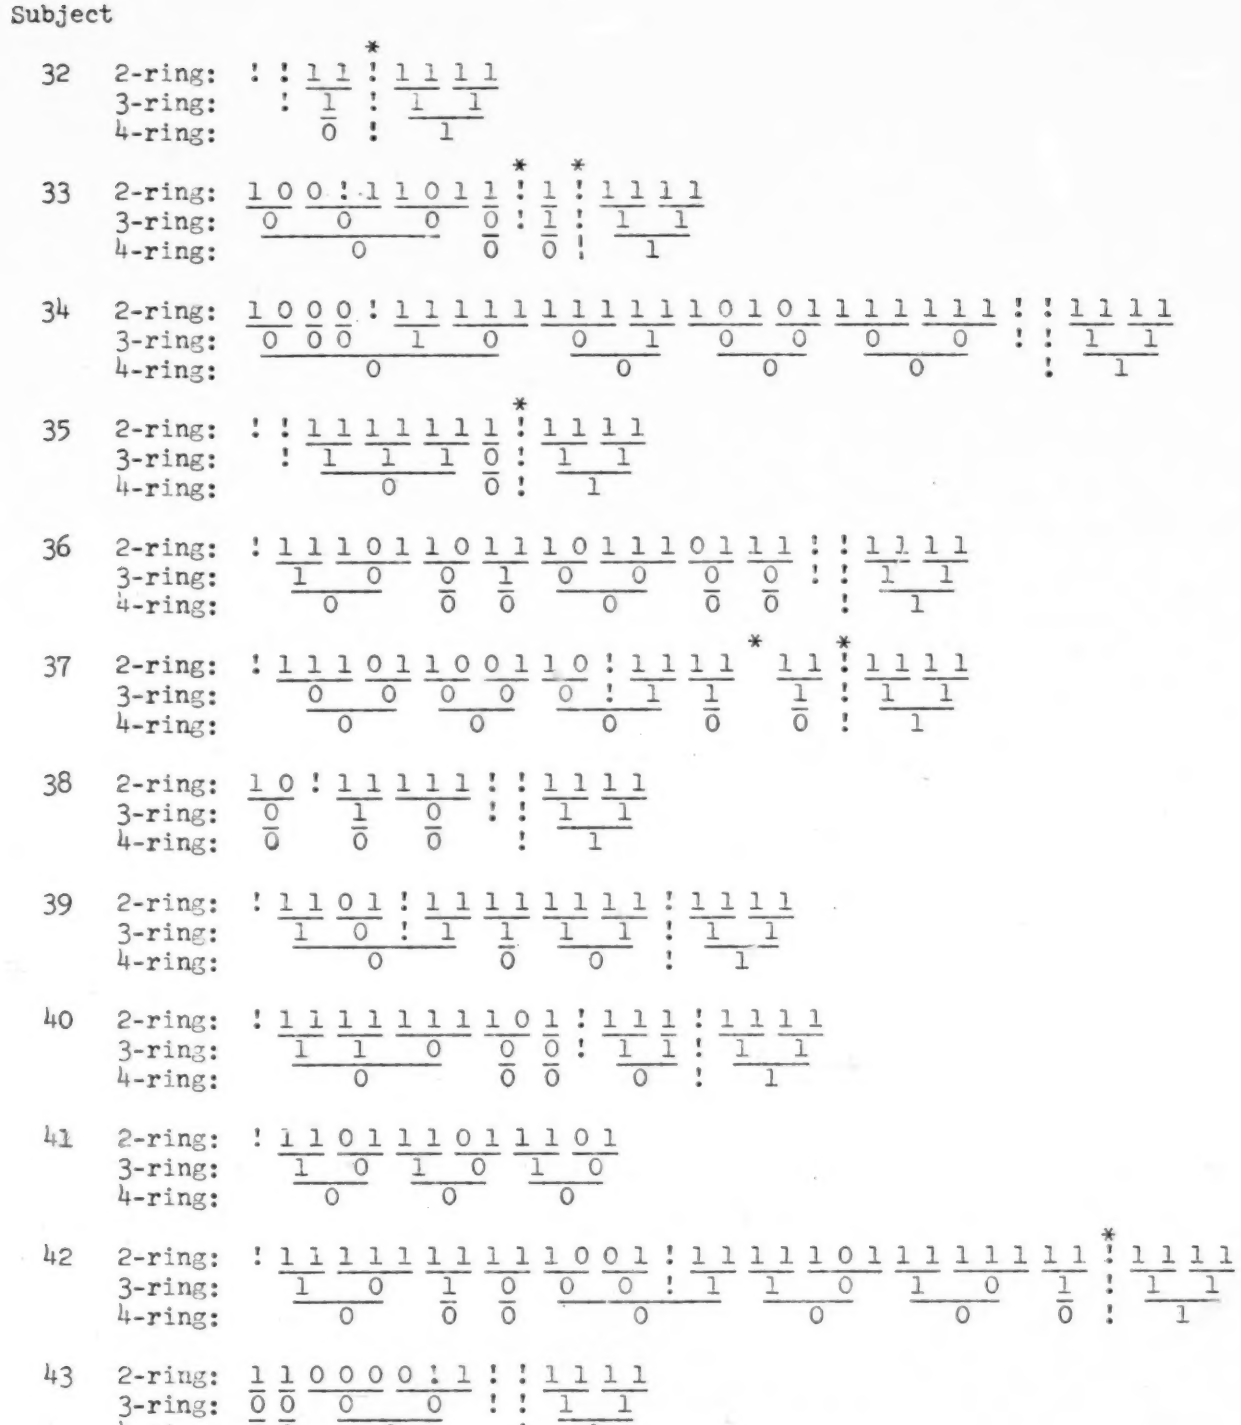

TABLE 4 (continued).

Subject

44	2-ring:	!	1	1	1	1	!	1	1	1	1	!	1	1	1	1			
	3-ring:		1		0		!	1		1			1		1				
	4-ring:				0								0						
45	2-ring:	!	!	1	1		!	1	1	1	1								
	3-ring:		!	1			!	1		1									
	4-ring:			0			!			1									

The foregoing tabulation of data does not include six subjects who solved the problem perfectly on their first attempt.